Handbook of Geospatial Intelligence

Centre for Administrative Services

Table of Contents

Chapter 1 — What Geospatial Intelligence Is....1

1.1 Definition and Origins....1

1.2 The Historical Evolution of the Discipline....2

1.3 What the Analyst Does....3

1.4 The Spatial Dimension....4

1.5 The Four Domains....4

1.6 Government, Commercial, and Academic Geospatial Work....5

1.7 Why Governments Fund Geospatial Intelligence....5

1.8 What the Office Produces....6

1.9 What Geospatial Intelligence Is Not....7

1.10 Worked Example: Producing a Flood Inundation Assessment....7

Chapter 2 — The Geospatial Intelligence Function in Government....9

2.1 Where the Function Sits....9

2.2 The Coordinating Character....10

2.3 Conflicts of Interest....11

2.4 The Revenue-Neutral Model....11

2.5 Governance Arrangements....12

2.6 The Political Economy of Placement....13

2.7 The Establishment Decision....13

2.8 Worked Example: Establishing a Geospatial Intelligence Function....14

Chapter 3 — Geospatial Intelligence and Adjacent Disciplines....15

3.1 Why Boundaries Matter....15

3.2 Surveying and Cadastre....16

3.3 Cartography....17

3.4 Spatial Planning....18

3.5 Statistics....18

3.6 Remote Sensing....19

3.7 The Scope Creep Pattern....20
3.8 Maintaining the Boundaries in Practice....21
3.9 Worked Example: Applying the Intake Test....22
Chapter 4 — The Geospatial Officer as a Professional Category....23
4.1 The Role Family....23
4.2 The Rank Structure....24
4.3 Career Paths....25
4.4 Certification and Qualifications....26
4.5 Language Requirements....27
4.6 Payscale Considerations....28
4.7 The Professional Identity....28
4.8 Worked Example: Recruiting and Developing a Hazard Specialist....29
Chapter 5 — The Commercial Firewall as Foundational Principle....30
5.1 The Principle....30
5.2 Comparative Models....31
5.3 The Statutory Basis....32
5.4 The Financial Architecture....33
5.5 The Audit Regime....33
5.6 The Stress Tests....34
5.7 The Firewall in Practice....35
5.8 Worked Example: A Commercial Commission Under Stress....35
Chapter 6 — The Acquisition-to-Archive Pipeline....36
6.1 The Pipeline Concept....37
6.2 The Six Capabilities....37
6.3 The Shared-Backbone Principle....39
6.4 The Swappable Modeling Heads....39
6.5 The Provenance Chain....41
6.6 The Cadence Design....41
6.7 The Pipeline and the Domains....42
6.8 Worked Example: A Day in the Pipeline....42
Chapter 7 — Acquisition and Ingest....44

7.1 The Acquisition Function.....44
7.2 Satellite Tasking.....45
7.3 Imagery Purchasing.....46
7.4 Sensor Feed Ingestion.....47
7.5 License Management.....47
7.6 The Provenance Record.....48
7.7 The Acquisition Discipline.....49
7.8 Worked Example: Acquiring Imagery for a Flood Event Response.....49
Chapter 8 — The Spatial Baseline.....51
8.1 The Base Layers.....51
8.2 The Digital Elevation Model in Detail.....52
8.3 The Cadastre Feed.....52
8.4 Base Layer Integrity.....53
8.5 The Synchronization Check.....54
8.6 The National Spatial Data Infrastructure.....55
8.7 The Governance of the Base Layers.....55
8.8 Worked Example: Resolving a Cadastre Discrepancy.....56
Chapter 9 — Processing.....57
9.1 The Processing Function.....57
9.2 Georeferencing.....58
9.3 Calibration and Mosaicking.....59
9.4 Projection Management.....60
9.5 Cloud Masking.....60
9.6 Data Fusion.....61
9.7 The Geometric Integrity Guarantee.....62
9.8 Worked Example: Processing Sentinel-2 Imagery for a Crop Condition Assessment.....62
Chapter 10 — Modeling.....63
10.1 The Modeling Function.....64
10.2 The Hydrology Plug-In.....64
10.3 The Crop Model Plug-In.....65
10.4 The Change Detection Plug-In.....66

10.5 The Prospectivity Plug-In....67
10.6 Model Validation....68
10.7 Model Provenance....69
10.8 Worked Example: A Coupled Hydrological-Hydraulic Model Run for a Flood Event....69
Chapter 11 — Production and Cartography....70
11.1 The Production Function....70
11.2 Cartographic Design....71
11.3 Symbology and Classification Marking....72
11.4 Format Conversion....73
11.5 The Publication Pipeline....73
11.6 The Cartographic Standards....75
11.7 Production and the Consumer....75
11.8 Worked Example: Producing a Flood Inundation Map for Event Response....76
Chapter 12 — Catalog and Archive....77
12.1 The Archive Function....77
12.2 The Catalog....78
12.3 Version History....78
12.4 Retention and Disposal....79
12.5 Access Control....80
12.6 The Audit Jurisdiction....80
12.7 The Archive and the Office....81
12.8 Worked Example: Archiving and Retrieving a Flood Inundation Map....82
Chapter 13 — The Office Manager....83
13.1 The Role of the Manager....83
13.2 The Manager's Daily Rhythm....84
13.3 The Morning Production Review....85
13.4 The Commercial Release Authorization....85
13.5 The Monthly Reporting Cycle....85
13.6 The Incident Response Activation....86
13.7 The Manager's Discipline....86
13.8 Worked Example: A Manager's Day During a Flood Event....87

Chapter 14 — The Technical Lead **88**
14.1 The Role of the Technical Lead 88
14.2 The Technical Lead's Daily Rhythm 89
14.3 The Daily Processing Queue Review 90
14.4 The Model Validation Sampling 90
14.5 The Base Layer Integrity Audit 91
14.6 The Specialist Support Consultation 91
14.7 The Technical Lead's Discipline 91
14.8 Worked Example: Diagnosing a Model Validation Failure 92
Chapter 15 — The Data Pipeline Operator **93**
15.1 The Role of the Data Pipeline Operator 93
15.2 The Operator's Daily Rhythm 94
15.3 The Morning Ingestion Check 95
15.4 The Processing Pipeline Configuration 96
15.5 The Weekly Base Layer Integrity Audit 96
15.6 The License Compliance Review 96
15.7 The Operator's Discipline 96
15.8 Worked Example: A Morning Ingestion Check with Multiple Issues 97
Chapter 16 — The Production Operator **98**
16.1 The Role of the Production Operator 99
16.2 The Operator's Daily Rhythm 99
16.3 The Product Production Run 101
16.4 The Classification Marking Verification 101
16.5 The Archival and Cataloguing 101
16.6 The Classified Access Log Review 101
16.7 The Operator's Discipline 102
16.8 Worked Example: Producing and Archiving a Commercial Prospectivity Report 103
Chapter 17 — The Hazard Specialist **104**
17.1 The Role of the Hazard Specialist 104
17.2 The Specialist's Daily Rhythm 105
17.3 The Flood Model Run 106

17.4 The Energy Infrastructure Monthly Monitoring Cycle....106
17.5 The Flood Event Response....107
17.6 The Change Detection Validation....107
17.7 The Specialist's Discipline....107
17.8 Worked Example: A Flood Event Response with Data Quality Issues....108
Chapter 18 — The Resource Economics Specialist....109
18.1 The Role of the Resource Economics Specialist....109
18.2 The Specialist's Daily Rhythm....110
18.3 The Weekly Crop Condition Bulletin....111
18.4 The Strategic Reserves Procurement Advisory....111
18.5 The Minerals Prospectivity Model Run....112
18.6 The Commercial Commission Production and Release....112
18.7 The Specialist's Discipline....112
18.8 Worked Example: A Week in the Life of the Resource Economics Specialist....113
Chapter 19 — Flood Risk Intelligence....114
19.1 The Domain Defined....114
19.2 The Scope....115
19.3 The Method....116
19.4 The Products....116
19.5 The Consumers....117
19.6 The Cadence....118
19.7 The Classification Regime....118
19.8 The Comparative Practice....119
19.9 Worked Example: Producing a 100-Year Flood Hazard Map for a Coastal City....120
Chapter 20 — Food Security Intelligence....121
20.1 The Domain Defined....121
20.2 The Scope....122
20.3 The Method....122
20.4 The Products....123
20.5 The Consumers....124
20.6 The Cadence....125

20.7 The Classification Regime 125
20.8 The Comparative Practice 126
20.9 Worked Example: Producing a Procurement Advisory During a Drought Year 127
Chapter 21 — Energy Infrastructure Intelligence 128
21.1 The Domain Defined 128
21.2 The Scope 128
21.3 The Method 129
21.4 The Products 130
21.5 The Consumers 131
21.6 The Cadence 131
21.7 The Classification Regime 132
21.8 The Comparative Practice 132
21.9 Worked Example: Detecting Vegetation Encroachment on a Transmission Corridor 133
Chapter 22 — Minerals and Critical Resources Intelligence 134
22.1 The Domain Defined 134
22.2 The Scope 135
22.3 The Method 136
22.4 The Products 137
22.5 The Consumers 138
22.6 The Cadence 138
22.7 The Classification Regime 139
22.8 The Comparative Practice 140
22.9 Worked Example: A Commercial Prospectivity Assessment for a Lithium Commission 140
Chapter 23 — The Cadastral Authority Interface 142
23.1 The Interface Defined 142
23.2 The Scope 142
23.3 The Method 143
23.4 The Products 144
23.5 The Consumers 145
23.6 The Cadence 146

23.7 The Classification Regime 147
23.8 The Comparative Practice 147
23.9 Worked Example: A Synchronization Discrepancy at the Riverbourne Basin 148
Chapter 24 — The Emergency Management Authority Interface 149
24.1 The Interface Defined 150
24.2 The Scope 150
24.3 The Method 151
24.4 The Products 152
24.5 The Consumers 153
24.6 The Cadence 154
24.7 The Classification Regime 155
24.8 The Comparative Practice 155
24.9 Worked Example: The Riverbourne Basin Flood Event Interface 156
Chapter 25 — The Strategic Reserve Authority Interface 158
25.1 The Interface Defined 158
25.2 The Scope 158
25.3 The Method 159
25.4 The Advisory Production 160
25.5 The Products 161
25.6 The Decision Feedback 162
25.7 The Classification Regime 163
25.8 The Comparative Practice 164
25.9 Worked Example: Producing the Heading Forecast Advisory During the Northfields Drought 165
Chapter 26 — The Sectoral Regulator Interface 166
26.1 The Interface Defined 166
26.2 The Scope 167
26.3 The Method 168
26.4 The Products 169
26.5 The Consumers 170
26.6 The Cadence 171

26.7 The Classification Regime 172

26.8 The Comparative Practice 172

26.9 Worked Example: The Highmoor-Portsmith Transmission Corridor Compliance Verification 174

Chapter 27 — The Water Authority Interface 175

27.1 The Interface Defined 175

27.2 The Scope 176

27.3 The Method 177

27.4 The Products 178

27.5 The Consumers 179

27.6 The Cadence 180

27.7 The Boundary Disputes 181

27.8 The Comparative Practice 181

27.9 Worked Example: The Riverbourne Basin Flood Event Coordination with the Water Authority 182

Chapter 28 — The National Statistics Agency Interface 184

28.1 The Interface Defined 184

28.2 The Scope 185

28.3 The Method 186

28.4 The Products 187

28.5 The Consumers 188

28.6 The Cadence 189

28.7 The Independence Protection 190

28.8 The Comparative Practice 190

28.9 Worked Example: The Portsmith Flood Exposure Assessment 191

Chapter 29 — The Geospatial Information System Specification 193

29.1 The Architectural Pattern 193

29.2 The Spatial Data Infrastructure 194

29.3 The Acquisition Desk 195

29.4 The Processing Engine 196

29.5 The Modeling Workbench 196

29.6 The Product Studio 197

29.7 The Catalog and Archive.....198
29.8 The Integration with the Core System.....199
29.9 Worked Example: The System Specification for the Riverbourne Basin Flood Event Response.....200
Chapter 30 — The Geospatial Role Family and Qualifications.....202
30.1 The Role Family.....202
30.2 The Rank Structure.....203
30.3 The Qualifications.....203
30.4 The Career Paths.....204
30.5 The Language Requirements.....205
30.6 The Training and Continuous Development.....206
30.7 The Payscale Considerations.....207
30.8 The Comparative Practice.....208
30.9 Worked Example: The Recruitment and Career Progression of a Hazard Specialist.....209
Chapter 31 — Training and the Administrative Training Authority.....210
31.1 The Training Pipeline.....210
31.2 The Foundational Course.....211
31.3 The Submodule Courses.....212
31.4 The Domain Courses.....213
31.5 The Leadership Courses.....214
31.6 The Continuing Education.....214
31.7 The Certification.....215
31.8 The Comparative Practice.....216
31.9 Worked Example: The Training Pipeline of a New Resource Economics Specialist.....217
Chapter 32 — Job Methods and Continuous Improvement.....218
32.1 The Job Methods Tradition.....219
32.2 The Method-Improvement Proposal.....219
32.3 The Annual Procedure Review.....220
32.4 The Post-Incident Review.....221
32.5 The Discipline of Not Freezing.....222
32.6 The Lean and Kaizen Traditions.....223

32.7 The Improvement Record....223
32.8 The Comparative Practice....224
32.9 Worked Example: The Post-Incident Review of the Riverbourne Basin Flood Event....225
Chapter 33 — Comparative Geospatial Intelligence Traditions....227
33.1 The Purpose of Comparative Study....227
33.2 The United States Geological Survey....227
33.3 The Copernicus Programme....228
33.4 The French Institut national de l'information géographique et forestière....229
33.5 Geoscience Australia....230
33.6 The Indonesian Badan Informasi Geospasial....231
33.7 The Canada Centre for Mapping and Earth Observation....231
33.8 The Lessons of Comparative Practice....232
33.9 Worked Example: A Comparative Analysis for a New Geospatial Office....233
Chapter 34 — The Geospatial Officer's Judgment....235
34.1 The Nature of Professional Judgment....235
34.2 When Models Disagree with Ground Truth....235
34.3 When Data Is Ambiguous....236
34.4 When the Firewall Is Tested....237
34.5 When the Event Demands More Than the Office Can Produce....238
34.6 When the Consumer Requests the Impossible....239
34.7 The Development of Judgment....240
34.8 The Judgment and the Rules....240
34.9 Worked Example: Exercising Judgment During the Riverbourne Basin Flood Event....241
Chapter 35 — Drafting and Communication for Geospatial Intelligence....243
35.1 The Variation by Domain....243
35.2 The Flood Bulletin....244
35.3 The Procurement Advisory....245
35.4 The Energy Infrastructure Risk Assessment....246
35.5 The Commercial Prospectivity Report....247
35.6 The Government Style Manual and the Plain Language Standard....248
35.7 The Cartographic Communication....248

35.8 The Briefing and the Oral Communication.....249

35.9 Worked Example: Drafting the Procurement Advisory for the Northfields Drought.....250

Chapter 36 — The Geospatial Officer's Toolkit.....251

36.1 The Toolkit's Purpose.....251

36.2 The Reference Works.....252

36.3 The Software.....253

36.4 The Data Sources.....254

36.5 The Professional Networks.....255

36.6 The Standards Bodies.....255

36.7 The Continuing Education and the Professional Development Path.....256

36.8 The Toolkit and the Handbook.....257

36.9 Worked Example: The Toolkit of Dr. Maren Holloway.....258

Appendix A — Classification and Retention Quick Reference.....259

A.1 The Classification Levels.....260

A.2 The Classification by Domain.....260

A.3 The Retention Schedule.....261

Appendix B — Satellite Sensor Specifications Quick Reference.....262

B.1 The Free-and-Open Optical Sensors.....262

B.2 The Free-and-Open Radar Sensors.....263

B.3 The Commercial High-Resolution Optical Sensors.....263

B.4 The Sensor Selection Guide.....263

Appendix C — Standards and Specifications Directory.....264

C.1 The Geographic Information Standards.....264

C.2 The Remote Sensing Standards.....265

C.3 The Domain-Specific Standards.....265

C.4 The Government Style Standards.....266

Appendix D — Comparative Institutions Directory.....266

D.1 The National Geospatial Agencies.....266

D.2 The Emergency Management Institutions.....267

D.3 The Water Management Institutions.....268

D.4 The Strategic Reserve Institutions.....268

D.5 The Training and Professional Institutions.....269

Appendix E — Worked Example Index and Cross-Chapter Continuity.....270

E.1 The Riverbourne Basin Flood Event Narrative.....270

E.2 The Northfields Drought Narrative.....271

E.3 The Highmoor-Portsmith Transmission Corridor Narrative.....271

E.4 The Officer Career Narratives.....272

Appendix F — Glossary and Acronyms.....272

F.1 The Glossary.....272

F.2 The Acronyms.....276

Index.....279

Note: This Table of Contents is generated via field codes. To ensure page number accuracy after editing, right-click the table and select "Update Field."

Chapter 1 — What Geospatial Intelligence Is

Part I — Foundations: The discipline of geospatial intelligence, its place in government, and the foundational principles that shape the office.

1.1 Definition and Origins

Geospatial intelligence is the production of analytical assessments about land, water, resources, and infrastructure from spatially referenced data. The discipline takes data that is tied to geographic coordinates — satellite imagery, elevation models, parcel boundaries, sensor readings — and produces assessments that inform government decisions: where flood water will go, how crops are developing, whether infrastructure is at risk, where mineral deposits are likely to be found.

The term "geospatial intelligence" originated in the United States national security community. The National Geospatial-Intelligence Agency (NGA), established in 2003 through the merger of the National Imagery and Mapping Agency (NIMA) and the Defense Mapping Agency's legacy functions, adopted the term to describe the integration of imagery analysis, imagery intelligence, and geospatial information into a unified discipline. The NGA defines geospatial intelligence as "the exploitation and analysis of imagery and geospatial information to describe, assess, and visually depict physical features and geographically referenced activities on the Earth."

The term has since spread to civilian government contexts, where it describes the same integration applied to civilian domains: hazard assessment, resource management, infrastructure protection, and environmental monitoring. This spread reflects a genuine convergence of practice. The methods developed for military target analysis — the overlay of imagery with elevation data, the detection of changes between acquisitions, the fusion of multiple data sources — are the same methods that produce flood inundation maps, crop condition assessments, and mineral prospectivity studies. The convergence is the reason the term has been adopted across the government geospatial community.

This Handbook uses the term in its civilian sense. The civilian geospatial intelligence office produces intelligence for civilian consumers — emergency management authorities, strategic reserve authorities, sectoral regulators, and the public — using the same analytical methods that the national security community uses for its consumers. The distinction between the civilian and the national security senses of the term is institutional, not methodological. The methods treated in this Handbook are applicable in both contexts.

1.2 The Historical Evolution of the Discipline

Geospatial intelligence draws on three older disciplines that developed independently and converged over the late twentieth century: photogrammetry, remote sensing, and geographic information systems. Each tradition has its own history, professional societies, and institutional homes, and the convergence into a single discipline is a recent development that remains incomplete.

Photogrammetry — the science of making reliable measurements from photographs — dates to the mid-nineteenth century. Aimé Laussedat, a French military officer, is generally credited with the first systematic use of photographs for surveying, beginning in 1851. Albrecht Meydenbauer, a German architect, independently developed similar techniques in 1858 for the documentation of

historic buildings. The first systematic military application came in the First World War, when aerial photography from reconnaissance aircraft was used to produce trench maps of the Western Front. The interwar period saw the development of analytical photogrammetry, which used mathematical rather than optical-mechanical methods to reconstruct three-dimensional positions from overlapping photographs. The analytical plotter, introduced in the 1950s, replaced the optical-mechanical instruments of the earlier era and made photogrammetric measurement more precise and more adaptable. The transition to digital photogrammetry in the 1980s, with the introduction of fully digital workstations, completed the transformation from optical to computational methods.

Remote sensing — the acquisition of data about the Earth's surface from aircraft and satellites — expanded rapidly after the Second World War. Aerial photography had been used for military reconnaissance during the war, and the postwar period saw the development of more sophisticated airborne sensors, including infrared and radar systems. The CORONA programme, launched by the United States in 1960, was the first satellite reconnaissance system, providing photographic imagery of the Soviet Union and other areas of interest. CORONA imagery was declassified in 1995, making approximately 860,000 satellite photographs available to the public and the scientific community. Landsat 1 (originally designated ERTS-1), launched in 1972, was the first civilian satellite designed for Earth observation. Landsat provided 80-metre multispectral imagery of the entire Earth's surface on a regular revisit cycle, and its data was distributed freely to users worldwide. The Landsat programme, now in its ninth generation, has provided the longest continuous record of the Earth's land surface from space.

The commercial high-resolution satellite era began with IKONOS, launched in 1999 by Space Imaging (now part of Maxar Technologies), which provided 1-metre panchromatic and 4-metre multispectral imagery — the first commercial satellite imagery with resolution comparable to classified military systems. Other commercial providers followed: QuickBird (2001, 0.6m), GeoEye-1 (2008, 0.41m), WorldView-3 (2014, 0.31m). The European Space Agency's Copernicus programme, beginning with the launch of Sentinel-1 in 2014, provided a new model: free and open Earth observation data from a dedicated satellite constellation. The Sentinel series (Sentinel-1 radar, Sentinel-2 optical, Sentinel-3 ocean and land, Sentinel-4 and -5 atmospheric) now provides the most comprehensive free Earth observation data source in the world.

Geographic information systems — software for storing, analyzing, and displaying spatial data — emerged in the 1960s. The Canada Geographic Information System (CGIS), developed by Roger Tomlinson in 1963 for the Canada Land Inventory, is generally considered the first operational GIS. CGIS was designed to overlay multiple datasets — soil, geology, land use, wildlife — to identify areas suitable for agriculture, forestry, or recreation, and the overlay capability it introduced remains the foundational operation of spatial analysis. ESRI (Environmental Systems Research Institute), founded by Jack Dangermond in 1969, launched ARC/INFO in 1982, which became the dominant commercial GIS software. GRASS (Geographic Resources Analysis Support System), developed by the US Army Corps of Engineers Construction Engineering Research Laboratory in 1982, provided an open-source alternative. QGIS (originally Quantum GIS), first released in 2002, has since become the most widely used open-source GIS desktop application.

The convergence of these three traditions was driven by two forces. The first was the digital revolution of the 1990s, which made it possible to integrate imagery, elevation data, and vector data

in a single analytical environment. Before the digital revolution, a photogrammetrist, a remote sensing specialist, and a cartographer worked with different tools on different products. After the digital revolution, they worked with the same data in the same software, and the boundaries between the traditions dissolved. The second force was the policy recognition that the government's geospatial functions were fragmented across agencies in ways that produced redundancy, inconsistency, and missed opportunities. The United States National Spatial Data Infrastructure policy, established by Executive Order 12906 in 1994, directed federal agencies to coordinate their spatial data activities and to make their data available to other agencies and the public. The European Union's INSPIRE Directive (Directive 2007/2/EC), adopted in 2007, established a comparable framework for European Union member states. These policy initiatives encouraged the convergence of the traditions into a unified discipline.

The convergence is not complete. National geospatial agencies still employ photogrammetrists, remote sensing specialists, and GIS analysts as distinct professional categories, and the professional societies for each tradition — the American Society for Photogrammetry and Remote Sensing (founded 1934), the International Society for Photogrammetry and Remote Sensing (founded 1910), the Urban and Regional Information Systems Association (founded 1963) — remain active. But the convergence has progressed far enough that the term "geospatial intelligence" is widely understood as describing an integrated discipline, and the role family treated in Chapter 4 reflects the integrated discipline rather than the separate traditions.

1.3 What the Analyst Does

Geospatial intelligence is distinguished from the disciplines that produce its inputs by what the analyst does. Surveying measures the location and elevation of points on the Earth's surface. Cartography transforms those measurements into maps. Remote sensing acquires imagery from satellites and aircraft. The analyst takes these outputs and produces assessments that support decisions. The analyst depends on the surveyor's measurements, the cartographer's representations, and the remote sensing specialist's imagery, and cannot work without them. But the analyst's product is different from the inputs, because the analyst adds the interpretive judgment that turns data into assessment.

Consider a flood inundation map. The map shows where water is, but the assessment behind it is the work of a hydrologist who has run a hydraulic model — a two-dimensional hydrodynamic model such as HEC-RAS 2D, MIKE 21, or TUFLOW — against a digital elevation model derived from LiDAR data, calibrated the model's roughness coefficients (Manning's n values) against observed channel conditions, set the boundary conditions from the rainfall and river stage data, and validated the model's output against observed flood extents from historical events. The map is the communication of the assessment; the assessment is the intelligence. The same distinction applies across domains. A crop condition bulletin is a map of vegetation greenness (the Normalized Difference Vegetation Index, NDVI, or the Enhanced Vegetation Index, EVI), but the assessment behind it is the work of an agricultural analyst who has interpreted the vegetation indices in light of the crop's growth stage, the weather conditions, the soil moisture, and the historical yield data. A mineral prospectivity map is a map of geological formations, but the assessment behind it is the work of a geologist who has integrated geological, geophysical (aeromagnetic, gravity), and geochemical

data through a prospectivity model — weights-of-evidence, fuzzy logic, or machine learning — to identify where mineral deposits are likely to be found.

National geospatial agencies illustrate the distinction in their institutional arrangements. The United States Geological Survey maintains the National Map — a cartographic product — and also produces the National Water Model, an analytical product that the Federal Emergency Management Agency uses for flood forecasting. The National Map is the input to the National Water Model; the National Water Model is the intelligence that the emergency management authority consumes. The French Institut national de l'information géographique et forestière maintains the cadastre and the topographic map series — cartographic products — and also produces flood risk assessments for the French emergency management authority. The cadastre and the topographic maps are the inputs; the flood risk assessments are the intelligence. The distinction between the cartographic product and the analytical product is the distinction between the mapping function and the intelligence function, and the analytical function is what makes the institution an intelligence producer rather than merely a mapping agency.

1.4 The Spatial Dimension

Geospatial intelligence uses geographic location as an analytical dimension. A security service report may mention a place name, an economic agency may reference an address, a diplomatic cable may cite a coordinate — but these are labels, not analytical frameworks. Geospatial intelligence uses location as the framework: it asks what the geographic position of a feature tells us about its characteristics, its risks, or its relationships to other features.

This dimension makes certain questions answerable that non-spatial methods cannot address. How many people live within the 100-year floodplain? Which parcels of agricultural land are showing crop stress? How close is vegetation encroaching on transmission corridors? Which geological formations have the highest prospectivity for critical minerals? Each of these questions requires overlaying, intersecting, or comparing spatial datasets — operations that only spatial methods can perform.

The spatial dimension is maintained through three components. The coordinate reference system is the mathematical framework that relates coordinates to positions on the Earth's surface. The most widely used global system is WGS 84 (World Geodetic System 1984), which is the reference system for the Global Positioning System. National and regional systems include NAD 83 (North American Datum 1983) in North America, ETRS 89 (European Terrestrial Reference System 1989) in Europe, and GDA 94 (Geocentric Datum of Australia 1994) in Australia. Geographic coordinates (latitude and longitude) are suitable for data storage and exchange but are not suitable for measurement, because the distance represented by one degree of longitude varies with latitude. Projected coordinate systems — such as the Universal Transverse Mercator (UTM), which divides the Earth into 60 zones, or national grid systems — transform geographic coordinates into planar coordinates where distance, area, and bearing can be measured. Every map projection introduces distortions in area, shape, distance, or direction, and the choice of projection is determined by which distortions are acceptable for the analysis at hand.

The base layers are the maintained datasets that provide the geographic context for analysis. The digital elevation model (DEM) represents the Earth's surface as a grid of elevation values. Global

DEMs include SRTM (Shuttle Radar Topography Mission, 30-metre resolution, 2000) and ASTER GDEM (Advanced Spaceborne Thermal Emission and Reflection Radiometer Global Digital Elevation Model, 30-metre resolution, 2009). Higher-resolution DEMs are produced from airborne LiDAR (Light Detection and Ranging), which can achieve 1-metre or sub-metre resolution. Land cover datasets classify the Earth's surface by cover type; global datasets include GlobeLand30 (30-metre, 2010) and ESA World Cover (10-metre, 2020), while regional datasets include the CORINE Land Cover (European Union, 100-metre) and the National Land Cover Database (United States, 30-metre). Hydrography datasets represent surface waters; examples include the National Hydrography Dataset (United States) and EU-Hydro (European Union). Administrative boundaries are maintained by national authorities and compiled globally by initiatives such as GADM (Database of Global Administrative Areas). The cadastre — the official, legally binding map of land boundaries and titles — is maintained by cadastral authorities and consumed by the geospatial office as a base layer.

The geocoding conventions are the rules by which non-spatial data is associated with spatial locations. Population statistics from a census are associated with census enumeration areas; agricultural reports are associated with parcels or administrative units; infrastructure inventories are associated with point locations or linear features. Geocoding is the mechanism that allows the office to integrate its non-spatial data with its spatial data, and without it, the office cannot produce the exposure assessments, the crop stratifications, or the infrastructure risk maps that constitute its intelligence.

1.5 The Four Domains

This Handbook treats four domains: flood risk, food security, energy infrastructure, and minerals and critical resources. These are not the only applications of geospatial intelligence — the discipline is also applied to disease surveillance, refugee tracking, maritime domain awareness, environmental compliance, and many other areas — but the four domains are the applications that a revenue-neutral geospatial office serving a civilian government is most likely to undertake, and together they exercise the full range of the discipline's methods, cadences, and classification regimes.

The four domains share the same pipeline, the same base layers, and the same information system. They share the same office, the same complement, and the same budget. This sharing is the reason for the pipeline architecture treated in Part II: the pipeline allows the four domains to share the non-analytical capabilities (acquisition, processing, production, archive) while maintaining domain-specific analytical capabilities through separate modeling plug-ins.

The four domains were chosen because they span the range of the discipline's civilian applications. Flood risk intelligence is a public-service function with event-driven production and a tiered classification regime. Food security intelligence is a public-service function with scheduled production and a split classification regime. Energy infrastructure intelligence is a public-service function with scheduled production and a security-confidential classification regime. Minerals intelligence is a commercial function with project-based production and a mixed classification regime. Together, they cover hazard assessment, resource management, infrastructure protection, and commercial intelligence — the principal civilian applications of the discipline.

1.6 Government, Commercial, and Academic Geospatial Work

Geospatial intelligence is produced in three sectors: government, commerce, and academia. The three sectors use similar methods but serve different purposes and operate under different accountability frameworks. The government office produces intelligence for government decisions, under government accountability, with the obligation to serve the public interest. A commercial firm produces geospatial analysis for its customers, under commercial accountability, with the obligation to serve its shareholders. An academic institution produces geospatial research for the scholarly community, under academic accountability, with the obligation to advance knowledge.

The government office operates under four institutional frameworks that shape its work. The classification framework governs how intelligence is marked, handled, and released; it is a matter of law, not office discretion. The procurement rules govern how data and services are acquired; they exist to ensure fair competition, value for money, and transparency. The audit regime governs how the office's work is examined and held to account; the national audit authority verifies that the office is doing what it claims to do. The records management standards govern how the office's outputs are preserved and made accessible; they are the condition of the office's institutional memory.

Officers who have worked in the commercial or academic sectors, where these frameworks are different or absent, sometimes perceive them as impediments. The classification framework ensures that sensitive intelligence is protected. The procurement rules ensure that public money is spent fairly. The audit regime ensures that the office is accountable. The records management standards ensure that the office's work is preserved. These are the conditions under which government intelligence is legitimate. An office that circumvents them produces intelligence that is not legitimate, however analytically sound it may be.

1.7 Why Governments Fund Geospatial Intelligence

Governments fund geospatial intelligence because it improves decisions, reduces economic losses, and protects public safety. These three categories — decision value, economic value, and public safety value — are the components of the case for the geospatial intelligence function.

The decision value is the value of better decisions. Flood risk intelligence informs evacuation decisions. Food security intelligence informs procurement decisions. Energy infrastructure intelligence informs regulatory decisions. Minerals intelligence informs stockpile decisions. Each of these decisions has consequences, and the intelligence that informs them has value to the extent that it improves the outcome. The value is not always quantifiable, because the counterfactual — what would the decision have been without the intelligence? — is not always knowable. But the value is real, and governments invest in geospatial intelligence because the investment produces better decisions than the government would make without it.

The economic value is the value of losses avoided and efficiencies gained. The commercial minerals arm sells prospectivity maps that reduce exploration risk, which lowers the cost of mineral discovery and production. The public-service arm produces flood risk intelligence that informs land-use planning, reducing property damage. It produces food security intelligence that informs reserve management, reducing price volatility. It produces energy infrastructure intelligence that informs infrastructure protection, reducing the cost of failures. The economic value of the United States

Geological Survey's earthquake hazard maps, for example, has been estimated in the billions of dollars through the informed building codes and land-use decisions they enable.

The public safety value is the value of protecting life and property. Flood intelligence that supports an evacuation order saves lives. Infrastructure intelligence that identifies a vulnerability before it is exploited prevents failures. Food security intelligence that triggers a timely procurement prevents shortages. This is the value the public most readily understands, and it is the value most often cited in the political justification for the office's budget.

1.8 What the Office Produces

The geospatial office delivers its intelligence in several product types, each suited to a different purpose and audience. The principal types are maps, bulletins, advisories, alerts, and reports.

Maps are the cartographic representation of the intelligence. The flood inundation map, the crop condition map, the infrastructure risk map, and the mineral prospectivity map are all maps. Maps are used in operational systems — the emergency management authority's mapping platform, the energy sector entity's asset management system — where the consumer needs to see the spatial pattern of the intelligence. A map shows where, but it does not easily convey the model assumptions, the data uncertainties, or the analytical qualifications behind the assessment. A map without accompanying analytical documentation may be misinterpreted.

Bulletins are periodic publications that summarize the office's intelligence for a domain. The weekly crop condition bulletin, the monthly infrastructure monitoring bulletin, and the seasonal flood outlook are all bulletins. Their regularity is their principal virtue: the consumer knows when to expect them, what they will contain, and how to use them. Their limitation is generality: a bulletin is written for a broad audience and may not contain the specific detail a particular consumer needs for a particular decision.

Advisories are targeted products that address a specific decision. The procurement advisory, the evacuation advisory, and the maintenance advisory are all advisories. An advisory addresses a specific question, provides a specific answer, and supports a specific decision. It is relevant only for the decision it addresses, and it loses its value if not acted upon promptly.

Alerts are urgent products that communicate a developing situation. The flood alert, the encroachment alert, and the infrastructure failure alert are all alerts. An alert is produced and delivered as quickly as the office can manage, because the consumer needs the information immediately. An alert cannot convey the full analysis behind it; it is followed by a fuller product when time allows.

Reports are comprehensive products that document the office's intelligence for the record. The flood risk assessment report, the annual mineral resource assessment, and the infrastructure vulnerability study are all reports. A report documents the analysis, the data, the methods, and the conclusions. It is retained for reference and planning. A report is not the product for a time-critical decision; producing a report when an alert or an advisory was needed is producing the wrong product.

1.9 What Geospatial Intelligence Is Not

Geospatial intelligence is not surveying, not cartography, not spatial planning, not statistics, and not remote sensing, though it draws on all five. The boundaries between geospatial intelligence and

these adjacent disciplines are treated in detail in Chapter 3. The boundaries matter because they determine which institutional functions belong in the geospatial office and which belong elsewhere. An office that absorbs the functions of its adjacent disciplines will find itself doing surveying, cartography, planning support, statistical analysis, or remote sensing processing instead of producing intelligence, and its analytical output will decline.

The boundaries are not always clear in practice. A request for a map of existing conditions, where no analysis is required, is a cartographic request that belongs in a mapping agency. A request for an assessment of flood exposure, where the map is the communication of an analytical result, is an intelligence request that belongs in the geospatial office. The office's intake function — which determines which requests the office accepts and which it declines — is the operational mechanism for maintaining the boundaries. The intake function is treated in Chapter 3, section 3.7.

1.10 Worked Example: Producing a Flood Inundation Assessment

The Riverbourne Basin, a fictional river basin of approximately 2,400 square kilometres, experiences 180 millimetres of rainfall over 48 hours. At 03:00 on the second day, the river stage sensors operated by the national hydrological service show the main river rising above the 5-year flood threshold at three of eight gauge stations. The office's event-driven production is activated by the office manager under the incident response activation procedure (section 13.6).

The data pipeline operator retrieves the latest sensor feed data: 15 rain gauges, 8 river stage sensors, and 3 weather stations. All feeds are active and delivering current data. The operator also retrieves the most recent satellite imagery (Sentinel-2, acquired at 06:00, 40 per cent cloud cover over the basin) and the LiDAR-derived digital elevation model (1-metre resolution, acquired in 2019, the most recent available). The operator ingests the data into the processing engine and records the provenance: data source, acquisition date, license terms, and intended domain consumer.

The hazard specialist configures the hydrology plug-in with event-specific parameters. The observed rainfall from the 15 rain gauges and the forecast rainfall from the national meteorological service are used as input to the rainfall-runoff model. The observed river stages from the 8 gauge stations are used as boundary conditions for the hydraulic model. The specialist runs a two-dimensional HEC-RAS model against the LiDAR DEM, producing an inundation extent and depth map at 5-metre grid resolution. The model run takes approximately 45 minutes on the office's processing engine.

The specialist validates the model run against the observed river stages at the 8 gauge locations. The model's predicted stages are within 15 centimetres of the observed stages at 7 of 8 gauges. The 8th gauge shows a 30-centimetre discrepancy, which the specialist attributes to a bridge obstruction that is not captured in the DEM. The specialist documents the discrepancy in the model provenance record, noting that the model may overestimate or underestimate the inundation extent in the vicinity of the bridge.

The specialist overlays the inundation map with the cadastre feed, identifying 347 parcels affected by the flooding: 23 residential, 4 commercial, and 320 agricultural. The specialist also overlays with the population data from the national statistics agency, estimating approximately 1,200 people within the inundated area, and with the infrastructure asset registry from the energy sector entity, identifying 2 substations and 1 transmission corridor within the inundated area.

The specialist generates the event-specific products. An inundation map is produced with Internal classification, because it includes critical infrastructure exposure (the substations and transmission corridor). An exposure report is produced with Internal classification for the same reason. Evacuation guidance is produced with Public classification, for public communication through the emergency management authority.

The production operator applies the cartographic template, the symbology, and the classification markings. The operator converts the products to the emergency management authority's required formats: GeoTIFF for the inundation map, PDF for the exposure report, and Shapefile for the evacuation routes. The operator submits the draft products to the office manager for review.

The office manager reviews the products, verifies the classification markings (Internal for the inundation map and exposure report, Public for the evacuation guidance), and confirms that the provenance record is complete. The manager authorizes release. The inundation map and exposure report are delivered through the secure data interface to the emergency management authority. The evacuation guidance is delivered through the public portal.

The specialist documents the model run in the catalog and archive submodule: model version (HEC-RAS 2D, version 6.3), parameters (event-specific rainfall, observed boundary conditions, 5-metre grid), validation results (7 of 8 gauges within 15 centimetres, 1 gauge with 30-centimetre discrepancy attributed to bridge obstruction), products generated (inundation map, exposure report, evacuation guidance), and the complete provenance chain from the sensor feed data through the processing, modeling, and production stages to the archive.

This worked example shows how the four elements of the definition — analytical purpose (the specialist's model run and validation), spatial dimension (the coordinate-referenced data and the cadastral overlay), domain scope (flood risk), and institutional character (the classification markings, the provenance chain, the audit trail) — manifest in the actual work of producing a flood inundation assessment.

Chapter 2 — The Geospatial Intelligence Function in Government

2.1 Where the Function Sits

The geospatial intelligence function is placed in government in one of three principal homes: under an interior or home affairs ministry, under an environment or natural resources ministry, or as a standalone agency. The choice determines the office's primary consumers, its peer agencies, and the conflicts of interest it must navigate. Each placement has comparative exemplars, and each produces a different office.

The interior ministry placement serves emergency management, border management, and domestic security as its primary consumers. The interior ministry is the coordinating ministry for internal affairs, and placing the geospatial office within it positions the office to serve the coordinating function. The interior ministry does not regulate the mineral sector, which avoids the conflict between the office's commercial minerals arm and the sectoral regulator of minerals — a conflict that arises when the office is placed under an environment ministry that regulates the environmental impacts of extraction.

The environment ministry placement serves environmental monitoring, resource management, and agricultural monitoring as its primary consumers. The environment ministry is the sectoral ministry for the natural environment, and placing the geospatial office within it positions the office to serve the environmental and resource functions. The placement creates a conflict of interest when the office operates a commercial minerals arm, because the environment ministry typically regulates the mineral sector, and the same ministry cannot both regulate the sector and sell it intelligence without a structural conflict.

The standalone agency placement serves a broad cross-government clientele and is established with institutional independence comparable to that of a national statistics agency. The standalone model is exemplified by the United States Geological Survey, which reports to the Department of the Interior but operates with substantial scientific autonomy; by Geoscience Australia, which reports to the Minister for Resources; and by the French Institut national de l'information géographique et forestière, which operates under the Ministry of Ecology.

The United States Geological Survey was established by the Organic Act of 1879 as a scientific agency within the Department of the Interior. The Survey's original mission was to classify public lands and examine the geological structure, mineral resources, and products of the national domain. Over the following decades, the Survey's mission expanded to include topographic mapping (transferred from the Army Corps of Engineers in 1882), water resources (the Water Resources Division was established in 1894), and later earthquake hazards, biological resources, and Earth observation. The Survey's placement within the Department of the Interior reflects its origins as a survey of public lands, which are managed by the Interior Department. The Survey now operates across hazard, resource, environmental, and mapping programmes, with its National Map programme providing the base cartographic data for the entire federal government.

Geoscience Australia was established in 2001 through the merger of the Australian Geological Survey Organisation (AGSO) and the Australian Surveying and Land Information Group (AUSLIG).

The merger consolidated the geological survey and the national mapping function in a single agency. The agency reports to the Minister for Resources and operates under the Geoscience Australia Act 2000. The agency's placement under the Minister for Resources reflects its origins as a geological survey and its continuing relationship with the mineral sector, which is a major contributor to the Australian economy. The agency's minerals division sells mineral intelligence to the exploration sector, and the revenue supports the agency's public geoscience programmes.

The French Institut national de l'information géographique et forestière (IGN) was established in 1940, evolving from the Service Géographique de l'Armée, the military mapping service. In 2012, IGN merged with the Inventaire forestier national (IFN, the national forest inventory), consolidating mapping, cadastre, and forest inventory in a single institution. IGN operates under the Ministry of Ecology (since 2009), a placement that reflects the French government's decision to integrate mapping, cadastre, and environmental monitoring. IGN maintains the cadastre, the topographic map series, the administrative boundaries, and the address database, and provides the integrated data to the government's analytical agencies through standardized interfaces.

The Ordnance Survey of Great Britain, established in 1791, is the world's oldest national mapping agency. The Ordnance Survey became a trading fund in 1999 under the Trading Funds Act 1973, a status that gave it commercial freedom while remaining in government ownership. The trading fund framework allows the Ordnance Survey to retain its revenue and to operate on a commercial basis, while requiring it to cover its costs from a combination of commercial revenue and government contracts. The Ordnance Survey reports to the Department for Business, Energy and Industrial Strategy.

Indonesia's Badan Informasi Geospasial (BIG) was established in 2011 as the national geospatial information agency, reporting directly to the President. BIG's placement reflects the priority given to geospatial information in a flood-prone archipelagic state of more than 17,000 islands, where the compound flood hazards — riverine flooding, coastal flooding, and storm surge — make geospatial intelligence a matter of national security.

2.2 The Coordinating Character

The geospatial intelligence function is a coordinating function, not a sectoral production function. A sectoral production function produces a specific good — food, energy, minerals, water — and sits under the ministry that regulates and promotes that sector. A coordinating function produces intelligence that supports decisions across multiple sectors, and sits where it can serve the coordinating and resilience functions of government.

This is why the geospatial office is not placed under the ministry of agriculture, the ministry of energy, or the ministry of mines. Each of those ministries is a sectoral producer. Placing the geospatial office under any one of them subordinates the coordinating intelligence function to a single sectoral interest. An office under the ministry of agriculture would produce food security intelligence that served the agricultural sector's interests. An office under the ministry of energy would produce energy infrastructure intelligence that served the energy sector's interests. The coordinating function — serving all sectors credibly — is lost.

The office serves multiple sectoral ministries, each with its own priorities and institutional culture. Its intelligence must be credible to all of them. The office's relationship with the emergency

management authority is different from its relationship with the strategic reserve authority, which is different again from its relationship with the sectoral regulator. The institutional placement must allow the office to maintain these relationships without being captured by any one of them.

The comparative practice confirms the coordinating character. The United States Geological Survey serves the Department of Homeland Security (for hazard assessment), the Department of Agriculture (for agricultural monitoring), the Department of Energy (for energy infrastructure), and the Department of Defense (for geospatial support), among others. None of these consumers is the Survey's parent department, and the Survey's placement within the Department of the Interior — which is not the parent of any of these consumers — allows the Survey to serve them all without being captured by any one. Geoscience Australia serves the Department of Climate Change, Energy, the Environment and Water, the Department of Defence, and the state geological surveys, among others. The coordinating character is the property that makes the office useful to all its consumers, and the institutional placement is the structural condition that preserves the coordinating character.

2.3 Conflicts of Interest

Each institutional placement creates or avoids specific conflicts of interest. The interior ministry placement avoids the conflict between the office's commercial minerals arm and the sectoral regulator of minerals, because the interior ministry does not regulate the mineral sector. The environment ministry placement creates that conflict, because the environment ministry typically regulates the environmental impacts of extraction. The same ministry cannot both regulate the mineral sector and sell it mineral intelligence without a structural conflict.

The conflict is not abstract. An office that sells mineral intelligence to the sector it regulates has a financial incentive to keep its data customers content, which is an incentive to soften its regulatory intelligence. If the office's budget depends on commercial minerals revenue, and the commercial minerals revenue depends on the goodwill of mining companies who are also regulated by the same ministry, the office faces pressure to produce intelligence that does not inconvenience its customers. The United States Geological Survey avoids this conflict by being placed in the Department of the Interior, which does not regulate mining on federal lands (that function belongs to the Bureau of Land Management, a separate bureau within the same department), and by maintaining the commercial firewall treated in Chapter 5.

A second conflict arises from the flood-response relationship. An office that produces flood risk intelligence for the same ministry that operates the flood response has an incentive to produce intelligence that justifies the response rather than intelligence that informs it. If the emergency management authority is a bureau within the same ministry as the geospatial office, and the ministry's leadership requests intelligence that supports a particular operational decision, the geospatial office faces pressure to produce that intelligence. The intelligence-response division, treated in Chapter 24, is the mechanism that prevents this pressure from distorting the intelligence, but the division is easier to maintain when the geospatial office and the emergency management authority are institutionally separate rather than subordinate to the same ministry.

These conflicts are the reason the firewall between the commercial and public-service arms, treated in Chapter 5, is a statutory arrangement rather than a managerial discretion. The conflict is structural, and it must be addressed structurally. The comparative practice of national geological

surveys that fund public geoscience through mineral-data sales — the United States Geological Survey, Geoscience Australia, the British Geological Survey — confirms that the structural separation is the only durable solution. Agencies that have relied on managerial discretion to maintain the separation have seen the separation erode under budgetary pressure, leadership turnover, and the gradual normalization of commingled work.

2.4 The Revenue-Neutral Model

The revenue-neutral model, in which the commercial arm of the geospatial office sustains the public-service arm, is a financing arrangement that shapes the office's institutional design. The model is exemplified by national geological surveys that sell mineral intelligence to the exploration sector and apply the revenue to the public geoscience functions conducted for the public benefit. It is also exemplified by national mapping agencies that sell high-resolution data and custom products to commercial customers and apply the revenue to the maintenance of the public base map.

The Ordnance Survey of Great Britain operates a commercial data business that funds a substantial portion of its public mapping function. Under the trading fund framework established by the Trading Funds Act 1973, the Ordnance Survey retains its revenue, operates on a commercial basis, and is required to cover its costs from a combination of commercial revenue and government contracts. The trading fund framework gives the Ordnance Survey commercial freedom — it can set prices, negotiate contracts, and invest in its business — while requiring it to maintain the public mapping function that is its statutory obligation. The framework includes a firewall: the Ordnance Survey's public sector mapping agreement provides free access to public sector users, while commercial users pay commercial rates, and the two sets of terms are separated by the trading fund's financial architecture.

The United States Geological Survey operates a hybrid model. Public data — Landsat imagery, National Map data, and most of the Survey's scientific publications — is freely available through the Earth Resources Observation and Science (EROS) Center. Value-added products and custom services are sold on a cost-recovery basis, and the revenue is deposited in the Survey's revolving fund. The revolving fund, authorized by statute, allows the Survey to accumulate revenue in good years and draw it down in poor years, so that the public-science programmes are not subject to the commercial cycle. The Government Accountability Office has examined the Survey's commercial activities and the revolving fund's management in multiple audits, and the audit findings have informed the Survey's ongoing refinement of its commercial practices.

Geoscience Australia's minerals division sells mineral intelligence to the exploration sector. The revenue supports the agency's public geoscience programmes, including the geological mapping, the mineral systems research, and the hazard assessment work that the agency conducts for the public benefit. The agency's enabling legislation, the Geoscience Australia Act 2000, establishes the conditions under which the agency may sell products and services and the conditions under which the revenue may be used.

The revenue-neutral model requires the commercial firewall discussed in Chapter 5, and the firewall is a consideration in the office's institutional placement. The placement must allow the commercial arm to operate with the pricing flexibility and customer orientation that cost-recovery requires, while protecting the public-service arm from the commercial pressure that would distort its

intelligence. The interior ministry placement, the environment ministry placement, and the standalone agency placement all accommodate the revenue-neutral model, but they accommodate it differently. The interior ministry placement places the commercial arm under the same financial regime as the public-service arm, which simplifies administration but requires careful firewall design. The standalone agency placement gives the commercial arm more institutional room, which eases the firewall but complicates the financial relationship with the treasury. The choice of placement and the design of the firewall are interdependent decisions.

2.5 Governance Arrangements

The governance arrangements for the geospatial office comprise the reporting line, the oversight mechanisms, and the inter-agency coordination frameworks. The reporting line runs from the office manager to the desk chief (or equivalent), to the bureau chief, to the deputy minister, and to the minister. The oversight mechanisms include the performance review by the desk chief (monthly), the bureau chief (quarterly), and the deputy minister (annually), and the audit by the national audit authority. The inter-agency coordination frameworks are the arrangements that govern the office's relationships with its consumers and partners, treated in Part V.

The governance arrangements are established by the office's enabling instrument — a statute, a regulation, or an organizational directive — which specifies the office's mandate, its reporting line, its complement, its budget, and its authorities. The enabling instrument is the legal foundation of the office, and the office that operates without a clear enabling instrument operates at the discretion of its parent ministry, which is a fragile foundation for a function that requires institutional independence.

The United States Geological Survey's authorities are established by the Organic Act of 1879 (as amended), which specifies the Survey's mandates and the conditions under which it may sell products and services. The Ordnance Survey's commercial activities are governed by the Ordnance Survey Act 1841 and the Trading Funds Act 1973, which establish the trading fund framework and the conditions that separate the commercial function from the public function. Geoscience Australia's commercial activities are governed by the Geoscience Australia Act 2000, which establishes the cost-recovery framework and the conditions that separate the commercial function from the public function. In each case, the enabling instrument provides the legal foundation that protects the office's independence and clarifies its authorities.

The office's complement, budget, and authorities are specified in the enabling instrument and reviewed periodically. The complement review determines whether the office's six seats are sufficient for the office's mandate, and whether additional seats should be added. The budget review determines whether the office's appropriated funding and commercial revenue are sufficient for the office's operations, and whether the balance between the two should be adjusted. The authorities review determines whether the office's authorities — the data acquisition authority, the inter-agency data-sharing authority, and the commercial release authority — are sufficient for the office's work, and whether additional authorities should be sought.

2.6 The Political Economy of Placement

The choice of institutional placement is not made in a vacuum. It is made in a political context where existing agencies have established interests, existing budgets have established claimants, and

existing reporting lines have established stakeholders. The geospatial function may be placed where it is because of historical accident, because of a ministerial reorganization, or because of a policy decision to consolidate fragmented functions. The placement that is analytically optimal may be politically infeasible, and the placement that is politically feasible may be analytically suboptimal.

The United States Geological Survey's placement under the Department of the Interior reflects a historical compromise. The Survey was established in 1879, before the Department of the Interior had assumed its current form. The Survey's placement within the Department was a political decision that preserved the Survey's scientific autonomy while giving it a departmental home. The decision was controversial at the time — the Survey's early directors, particularly John Wesley Powell, fought political battles with Congress over the Survey's budget and scope — and the settlement that emerged gave the Survey a degree of scientific independence that has been its institutional hallmark ever since.

Geoscience Australia's placement under the Minister for Resources reflects the agency's origins as a geological survey and its continuing relationship with the mineral sector, which is a major contributor to the Australian economy. The placement positions the agency to serve the mineral sector's geoscience needs, but it also means that the agency's public-science programmes — its hazard assessment, its environmental monitoring, its water resources work — are housed in a ministry whose primary constituency is the resource sector. The tension between the resource-sector orientation and the broader public-science mandate is managed through the agency's enabling legislation, which specifies the agency's full range of functions, and through the agency's governance arrangements, which include reporting to multiple ministerial portfolios.

The French IGN's placement under the Ministry of Ecology reflects the French government's decision to integrate mapping, cadastre, and environmental monitoring under a single ministry. The decision was part of a broader reorganization of the French government's environmental functions, and it positioned IGN to serve the environmental monitoring and land management functions that are the Ministry of Ecology's responsibilities. The placement changed IGN's institutional culture: from a military-origin mapping agency serving a broad government clientele, IGN became an environmental ministry agency serving the ministry's priorities, and the adjustment required a renegotiation of IGN's relationships with its other consumers.

The political economy of placement has practical consequences. A geospatial office placed under a ministry that does not understand its function may be starved of resources or diverted to the ministry's sectoral priorities. A geospatial office placed under a ministry that understands its function but has a competing interest may find its intelligence distorted by the ministry's policy preferences. A geospatial office placed as a standalone agency may find itself without a ministerial champion when budget cuts are proposed. The choice of placement should account for these political realities, not only for the analytical considerations.

2.7 The Establishment Decision

The establishment of a geospatial intelligence function is a decision that should be made on the basis of a needs assessment, a comparative review, and a political assessment. The needs assessment identifies the government's geospatial intelligence requirements and the gaps in the current arrangements. The comparative review examines how other jurisdictions have organized the function

and what lessons their experience offers. The political assessment identifies the feasible placements, given the existing institutional landscape and the political constraints.

The establishment decision should specify the office's mandate, its placement, its complement, its budget, its authorities, and its reporting line. The decision should be documented in an enabling instrument — a statute, a regulation, or an organizational directive — that provides the legal foundation for the office. The decision should be reviewed periodically, because the office's mandate and placement may need to evolve as the government's geospatial intelligence requirements evolve and as the institutional landscape changes.

The establishment decision is the most consequential institutional decision in the geospatial intelligence function's life. A well-made decision produces an office that serves its consumers credibly and sustainably. A poorly-made decision produces an office that is captured by a sectoral interest, starved of resources, or diverted to unrelated functions. The decision should be made deliberately, with the needs assessment, the comparative review, and the political assessment all informing the choice.

2.8 Worked Example: Establishing a Geospatial Intelligence Function

A jurisdiction with a population of approximately 15 million has its geospatial functions fragmented across multiple agencies. The national mapping agency maintains the base map but does not produce analytical assessments. The geological survey maintains the mineral data but does not commercially exploit it. The meteorological service maintains the weather data but does not produce flood risk assessments. The emergency management authority produces hazard assessments but uses outdated base maps. The agricultural ministry monitors crops but uses different satellite imagery than the hazard assessments. The statistical agency maintains the census data but does not share it with the other agencies in a standardized format. There is no coordinating geospatial intelligence function, and the agencies' products are inconsistent and incompatible.

A ministerial review identifies the gaps. The review finds that flood risk assessments produced by the emergency management authority use base maps that are three years out of date. Crop condition assessments produced by the agricultural ministry use different satellite imagery than the flood assessments, and the two cannot be cross-referenced. Infrastructure monitoring is not conducted at all, because no agency has the mandate or the capability. Mineral prospectivity data produced by the geological survey is not commercially exploited, because the survey has no commercial function and no authority to sell products.

The review commissions a needs assessment, which identifies four priority domains: flood risk (the jurisdiction has significant riverine and coastal flood hazards), food security (the jurisdiction maintains a strategic grain reserve), energy infrastructure (the jurisdiction has a national energy grid that is not monitored for encroachment), and minerals (the jurisdiction has significant mineral resources that are not commercially assessed). The needs assessment estimates the annual value of the geospatial intelligence function at approximately $50 million in losses avoided (flood damage reduction, food price volatility reduction, infrastructure failure prevention) and $10 million in commercial revenue (mineral intelligence sales to the exploration sector).

The review commissions a comparative review, which examines the United States Geological Survey, Geoscience Australia, and the French IGN. The review finds that all three agencies maintain

shared-backbone pipelines that serve multiple domains from a single acquisition and processing infrastructure. All three operate under enabling legislation that specifies their mandates, authorities, and commercial frameworks. The two that sell commercial products (the USGS and Geoscience Australia) maintain commercial firewalls that separate the commercial arm from the public-service arm through statutory mechanisms.

The review commissions a political assessment, which identifies three feasible placements. The interior ministry placement is feasible because the interior ministry already houses the emergency management authority (the office's principal consumer) and does not regulate the mineral sector (avoiding the conflict of interest). The environment ministry placement is feasible because the environment ministry already houses the geological survey (which would be merged into the new office), but it creates the minerals-regulation conflict. The standalone agency placement is feasible but lacks a ministerial champion, because no existing ministry is willing to cede the geospatial function to an independent agency.

The political assessment recommends the interior ministry placement. The interior ministry does not regulate the mineral sector, which avoids the conflict of interest. The interior ministry already houses the emergency management authority, which simplifies the intelligence-response coordination. The interior ministry has a minister who is willing to champion the new office, which provides the political support the office needs in its establishment phase.

The establishment decision specifies the office's mandate: geospatial intelligence across four domains (flood risk, food security, energy infrastructure, minerals). Its placement: a bureau within the interior ministry, with a dedicated desk (the Services Desk) reporting to the bureau chief of civil affairs. Its complement: six seats (office manager, technical lead, data pipeline operator, production operator, hazard specialist, resource economics specialist). Its budget: $8 million annually, of which $3 million from commercial revenue and $5 million from appropriated funding, with a revolving fund to smooth the commercial revenue. Its authorities: commercial release authority (sole, non-delegable), data acquisition authority (under the government's framework agreements), and inter-agency data-sharing authority (under the inter-agency data-sharing framework). Its reporting line: from the office manager to the desk chief, to the bureau chief, to the deputy minister of the interior.

The enabling instrument is drafted as a regulation under the interior ministry's statutory authority. The regulation specifies the four firewall mechanisms: separate charge codes for commercial and public-service work, separate classification markings for commercial and public-service products, sole commercial release authority vested in the office manager, and annual audit by the national audit authority. The regulation also specifies the office's mandate, its complement, its budget structure, its authorities, and its reporting line, and it establishes the office as a permanent operational unit within the interior ministry's organizational structure.

Chapter 3 — Geospatial Intelligence and Adjacent Disciplines

3.1 Why Boundaries Matter

Geospatial intelligence draws on five adjacent disciplines: surveying and cadastre, cartography, spatial planning, statistics, and remote sensing. Each discipline contributes data, methods, or institutional capacity that the geospatial office needs. But the geospatial office is not a surveying agency, a mapping agency, a planning authority, a statistical agency, or a remote sensing data centre. The boundaries between geospatial intelligence and these adjacent disciplines determine which institutional functions belong in the geospatial office and which belong elsewhere.

The boundaries matter for a practical reason. An office that absorbs the functions of its adjacent disciplines will find itself doing surveying, cartography, planning support, statistical analysis, or remote sensing processing instead of producing intelligence. Each absorbed function consumes staff time, budget, and managerial attention that should be directed at analytical work. The result is an office that does many things adequately but does its core function — producing analytical assessments — poorly or not at all. This pattern has been observed in government geospatial agencies across multiple jurisdictions: an office established to produce intelligence gradually becomes a mapping agency, a data-processing facility, or a planning support unit, and the analytical capability that justified its establishment atrophies.

The boundaries are maintained through institutional arrangements, not through territorial claims. A geospatial office that tries to maintain its boundaries by asserting that certain tasks are "not its job" will create friction with its consumers and partners. A geospatial office that maintains its boundaries by having clear intake criteria — documented, transparent, and applied consistently — will maintain its analytical focus without creating friction. The intake function, treated in section 3.8, is the operational mechanism for maintaining the boundaries.

The comparative practice of national geospatial agencies shows that offices with clear discipline boundaries produce better intelligence than offices with blurred boundaries. The United States Geological Survey's Organic Act of 1879 established the Survey as a scientific agency, not a mapping agency, and the Survey has maintained its analytical character throughout its history despite the expansion of its mapping functions. Geoscience Australia's enabling legislation specifies the agency's analytical functions separately from its data production functions. The French IGN, which has absorbed both the cadastre and the forest inventory, maintains an internal separation between its data production divisions and its analytical services, because the institutional separation within the agency is the mechanism that preserves the analytical character.

3.2 Surveying and Cadastre

Surveying measures the location and elevation of points on the Earth's surface. The surveyor uses instruments — total stations, GNSS receivers, levels, LiDAR scanners — to determine coordinates with a specified precision and accuracy. Cadastral surveying, the branch of surveying that concerns the geospatial office most directly, establishes and re-establishes the boundaries of land parcels. The cadastral authority maintains the cadastre: the official, legally binding map of land boundaries and titles. The cadastre is the record on which property rights, transactions, and disputes

are decided. It is maintained under a statutory framework, updated through legal instruments (deeds, subdivisions, consolidations), and defensible in court.

The geospatial office consumes the cadastre as a base layer. The cadastre is what makes the office's products actionable: a flood inundation map is meaningful when it can be associated with specific parcels; a crop condition assessment is meaningful when it can be associated with specific fields; an infrastructure risk assessment is meaningful when it can be associated with specific properties; a mineral prospectivity study is meaningful when it can be associated with specific exploration licenses. Without the cadastre, the office's intelligence is a map of where something is happening but not whose interests are affected.

The geospatial office does not survey, does not maintain the cadastre, and does not determine land boundaries. These functions belong to the cadastral authority, which has the legal mandate, the statutory framework, and the professional capacity to perform them. The geospatial office that surveys parcels or determines boundaries has exceeded its mandate and compromised the cadastre's legal authority, because the office's determinations are not made under the cadastral framework and are not defensible in court in the same way the cadastral authority's determinations are.

The institutional arrangements for the cadastre vary across jurisdictions. The French Institut national de l'information géographique et forestière maintains both the cadastre and the topographic map series, but the geological survey (Bureau de Recherches Géologiques et Minières), the agricultural monitoring service, and the flood risk assessment service are separate institutions that consume IGN's base data. The Singapore Land Authority maintains the cadastre and the national map, while the national water agency (PUB) and the National Environment Agency produce the flood risk and environmental intelligence that depend on it. The United States Geological Survey maintains the National Map but does not maintain the cadastre, which is a county-level function in the United States; the Survey's analytical work on hazards, resources, and environment consumes the cadastral data as an input, and the Survey does not modify the data.

The Netherlands' Cadastre, Land Registry and Mapping Agency (Kadaster) maintains the cadastre, the land registry, and the national topographic map. Kadaster was established in 1838 and operates as a non-departmental public body under the Ministry of the Interior and Kingdom Relations. The agency's integrated structure — cadastre, registry, and mapping in one institution — is a model that several jurisdictions have studied, but the integration does not extend to analytical work: the Dutch geological survey (TNO-Geological Survey of the Netherlands), the agricultural monitoring service, and the flood risk assessment service are separate institutions that consume Kadaster's base data.

The two functions — cadastral recording and geospatial analysis — have different integrity requirements. The cadastre must be legally authoritative. It is maintained under a statutory framework, updated through legal instruments, and defensible in court. A boundary determination in the cadastre has the force of law, and the cadastral authority's surveyors are licensed professionals whose work is subject to legal scrutiny. The geospatial intelligence product must be analytically sound. It is validated against ground truth, reproducible from its inputs, and honest about its uncertainties. A flood inundation map is an analytical assessment, not a legal determination, and its authority comes from the quality of the model and the validation, not from a statutory framework.

Combining the two functions in a single office creates pressure to make the cadastre serve the analysis, or to make the analysis serve the cadastre. If the same office maintains the cadastre and produces flood risk assessments, and the flood risk assessment reveals that a parcel boundary is wrong, the office faces a conflict: correcting the cadastre to match the assessment exceeds the analytical office's mandate, but failing to correct the cadastre means the assessment is based on incorrect base data. The geospatial office that respects the boundary consumes the cadastre as it is, reports discrepancies to the cadastral authority through the formal discrepancy reporting channel (treated in Chapter 23), and waits for the cadastral authority to resolve them before adjusting its products.

The discrepancy reporting is a two-way relationship. The cadastral authority maintains the cadastre, but the geospatial office, through its analytical work, may discover discrepancies that the cadastral authority's own processes have not identified. A flood model may reveal a parcel boundary that does not match the observed terrain. A crop condition assessment may reveal a parcel registered as agricultural but used for another purpose. An infrastructure monitoring run may reveal a parcel boundary inconsistent with the built infrastructure. The geospatial office reports these discrepancies to the cadastral authority, the cadastral authority investigates and resolves them through its statutory process, and the resolution is reflected in the next cadastre feed update. This two-way relationship is the constructive expression of the boundary: the cadastral authority provides the cadastre, and the geospatial office provides the discrepancy reporting that helps the cadastre stay current.

3.3 Cartography

Cartography transforms coordinate data and attribute data into maps. It develops the symbology, generalization, and design conventions that make spatial information legible. A well-designed map communicates the spatial pattern of the intelligence at a glance; a poorly designed map obscures the pattern and may mislead the consumer. The geospatial office produces maps as one of its product types, and the maps must meet cartographic standards. But the map is the communication of the intelligence, not the intelligence itself. The intelligence is the analytical assessment; the map is how the assessment is communicated.

The distinction matters because the cartographic standard and the analytical standard are different, and they are enforced differently. The cartographic standard concerns the map's symbology (are the symbols clear, consistent, and distinguishable?), its generalization (are features simplified appropriately for the map scale?), its layout (is the legend, scale bar, north arrow, and title block in the standard location?), and its typography (are the labels legible and correctly placed?). The analytical standard concerns the model's accuracy (does the inundation extent match the observed flood?), its completeness (are all affected parcels included in the exposure report?), and its honesty (are the model's uncertainties documented?). A map that is cartographically excellent but analytically wrong is a failure; a map that is cartographically rough but analytically sound may be a success, depending on the consumer and the purpose.

The United States Geological Survey's publication process separates the scientific review of its analytical products from the cartographic review of its published maps. The scientific review is conducted by peer scientists who evaluate the analytical methods, the data, the validation, and the conclusions. The cartographic review is conducted by cartographers who evaluate the symbology, the

generalization, the layout, and the typography. The two reviews are conducted by different teams with different criteria, and a product must pass both reviews before publication. This separation reflects the distinction between the analytical function and the cartographic function, and it prevents the cartographic standard from overriding the analytical standard (or vice versa).

The Ordnance Survey of Great Britain maintains a similar separation. The agency's cartographic standards are documented in the Ordnance Survey Map Specification, which specifies the symbology, generalization, and layout for every map product the agency publishes. The analytical work that uses Ordnance Survey data — flood risk assessments produced by the Environment Agency, agricultural monitoring produced by the Department for Environment, Food and Rural Affairs, infrastructure monitoring produced by the National Infrastructure Commission — is conducted by separate institutions that consume the Ordnance Survey's base data. The separation ensures that the Ordnance Survey's cartographic standards do not constrain the analytical work, and that the analytical work does not compromise the Ordnance Survey's cartographic standards.

The geospatial office's production function, treated in Chapter 11, includes cartographic production as a capability. The production operator applies the cartographic template, the symbology, and the layout to the specialist's analytical output. The production operator's work is cartographic, not analytical, and the production operator does not modify the analytical content. The office manager reviews the product for both cartographic compliance and analytical correctness before authorizing release. The separation of the cartographic function (production operator) from the analytical function (specialist) within the office is the internal expression of the boundary between cartography and geospatial intelligence.

3.4 Spatial Planning

Spatial planning organizes land use, settlement patterns, and infrastructure across geographic space. The spatial planner allocates: where growth will be encouraged, where it will be constrained, where infrastructure will be provided, where environmental assets will be protected. The geospatial office produces the intelligence that informs those decisions, but it does not make them. The distinction is the distinction between analysis and decision, between intelligence and policy, and it is maintained to protect the analytical independence of the geospatial office from the political pressures that bear on the planning function.

The Dutch national spatial strategy, produced under the Ministry of Infrastructure and Water Management, makes the strategic land-use decisions for the Netherlands. The strategy is revised approximately every ten years and is implemented through provincial structure plans and municipal zoning plans. The geospatial intelligence that informs the strategy is produced by a range of institutions: the Netherlands' Cadastre, Land Registry and Mapping Agency (Kadaster) provides the base data, the Royal Netherlands Meteorological Institute (KNMI) provides the climate data, the Deltares research institute provides the flood risk modeling, and the Netherlands Environmental Assessment Agency (PBL) provides the integrated assessments. None of these institutions makes the planning decisions; they produce the intelligence that the planning authority consumes.

The Singapore Concept Plan, produced by the Urban Redevelopment Authority (URA), makes the strategic land-use decisions for Singapore. The Concept Plan is revised approximately every ten years and is implemented through the Master Plan, which is revised every five years. The geospatial

intelligence that informs the plan is produced by the Singapore Land Authority (base data), the national water agency PUB (flood risk), the National Environment Agency (environmental monitoring), and other agencies. The URA consumes the intelligence and makes the planning decisions; the intelligence-producing agencies do not participate in the decisions.

The French system provides a further illustration. The Schémas de cohérence territoriale (SCoT) and the Plans locaux d'urbanisme (PLU) are the principal spatial planning instruments. They are produced by the communes and the intercommunal bodies, with the regional prefect providing oversight. The geospatial intelligence that informs the plans is produced by IGN (base data), the BRGM (geological hazard data), and the regional directorates of the central government ministries. The planning decisions are made by the communes; the intelligence is produced by the agencies. The separation is institutional: the communes do not produce the intelligence, and the agencies do not make the planning decisions.

The planning function operates on a political horizon. Planning decisions are made by elected officials or their appointees, and the decisions are subject to political contestation. The intelligence function operates on a technical horizon. Intelligence products are produced by professional analysts, and the products are subject to technical review. The planner may be tempted to request intelligence that supports a preferred policy outcome, and the intelligence office that reports to the planner may be tempted to produce it. This is why institutional placement, treated in Chapter 2, matters: a geospatial office that reports to a planning authority, or that is placed under a planning ministry, faces this pressure on every product. An office that maintains the distinction produces intelligence that is credible to all parties in the planning process, including the parties that disagree with the planner's preferred outcome. An office that does not maintain the distinction produces intelligence that is credible only to the planner, which is not intelligence at all.

The pressure to produce supportive intelligence is not always explicit. It may take the form of a request for "additional analysis" that happens to focus on the factors favouring the preferred outcome. It may take the form of a request to "present the results differently" that happens to emphasize the favourable aspects. It may take the form of a budgetary pressure that rewards compliant intelligence and penalizes inconvenient intelligence. The geospatial office that recognizes these pressures for what they are — attempts to compromise the analytical independence — and resists them through the discipline boundaries and the institutional placement is the office that produces credible intelligence. The office that does not recognize the pressures, or that recognizes them but yields, produces intelligence that serves the planner rather than the public.

3.5 Statistics

The national statistics agency collects, processes, and publishes quantitative data about a population, an economy, or a society. It produces the census, the economic indicators, and the social statistics. The geospatial office produces the spatial intelligence that informs, and is informed by, the statistical record. The two offices exchange data under the inter-agency data-sharing framework.

The United States Census Bureau, established as a permanent agency in 1902, produces the decennial census (mandated by Article I, Section 2 of the Constitution) and the American Community Survey (a continuous survey that replaced the long-form census in 2005). The Bureau's data is georeferenced to census tracts, block groups, and blocks, and is the primary source of population and

demographic data for the United States. The United States Geological Survey produces the spatial intelligence that informs and is informed by the census data: the Survey's hazard exposure assessments overlay flood, earthquake, and landslide hazard maps with the census population data to estimate the population at risk. The two agencies are institutionally separate, and the separation protects the Census Bureau's independence from the executive (established by Title 13 of the United States Code, which protects the confidentiality of census data and the independence of the Bureau's statistical work) and the Survey's independence from the sectoral ministries.

The Australian Bureau of Statistics, established as the Commonwealth Bureau of Census and Statistics in 1905, produces the census (conducted every five years) and a range of economic and social statistics. Geoscience Australia produces the spatial intelligence that informs and is informed by the census data. The two agencies are institutionally separate, and the separation is protected by the Census and Statistics Act 1905, which establishes the Bureau's independence.

Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS), established in 1899, produces the Dutch statistical record. The Netherlands' Cadastre, Land Registry and Mapping Agency (Kadaster) produces the spatial intelligence. The two agencies are institutionally separate, and the separation is protected by the Statistics Netherlands Act 1996, which establishes CBS's independence.

The separation protects the independence of both functions. The statistics agency must be independent of the executive to produce statistics that are trusted by the public and the legislature. The geospatial office must be independent of the sectoral ministries to produce intelligence that is trusted by all of them. Combining the two functions in a single office would create a conflict between the executive's interest in favorable statistics and the sectoral ministries' interest in favorable intelligence, and the conflict would compromise both. The United Nations Fundamental Principles of Official Statistics, adopted in 1994 and endorsed by the General Assembly, establish the independence of official statistics as a principle of international law, and the principle applies with equal force to the spatial intelligence that informs and is informed by the statistical record.

The data exchange between the statistics agency and the geospatial office is governed by the inter-agency data-sharing framework, treated in Chapter 28. The statistics agency provides the census data and the survey data to the geospatial office for use in the office's exposure assessments and model stratification. The geospatial office provides the spatial products — land cover data, infrastructure data, administrative boundaries — to the statistics agency for use in the agency's survey design and publication. The data exchange is conducted under the statistical confidentiality framework, which protects the personally identifiable information in the census and survey data and the commercially sensitive information in the infrastructure data. The confidentiality framework is the condition of the data exchange: without it, the statistics agency cannot share the data, and the geospatial office's exposure assessments lose their demographic foundation.

3.6 Remote Sensing

Remote sensing acquires imagery and other data from satellites, aircraft, and drones, and processes it into analysis-ready products. The remote sensing specialist tasks satellites, purchases imagery, ingests sensor feeds, and processes the data through georeferencing, calibration, mosaicking, projection management, cloud masking, and data fusion. The geospatial office acquires

and processes remote sensing data as part of its pipeline, but the acquisition and processing are capabilities in service of the analytical function, not the office's purpose.

The European Space Agency's Copernicus programme, established in 2014, is the world's largest free and open Earth observation data source. The programme's Sentinel satellite constellation comprises Sentinel-1 (C-band synthetic aperture radar), Sentinel-2 (multispectral optical), Sentinel-3 (ocean and land colour, sea surface temperature, altimetry), Sentinel-4 (atmospheric composition, geostationary), and Sentinel-5 (atmospheric composition, polar orbiting). The Copernicus Ground Segment acquires, processes, and distributes the Sentinel data freely to all users. The intelligence produced from the data is produced by the user institutions: national geological surveys use Sentinel-1 radar data for ground subsidence monitoring, national agricultural monitoring services use Sentinel-2 optical data for crop condition assessment, national flood forecasting services use Sentinel-1 radar data for flood extent mapping (radar penetrates clouds, which is critical for flood events that are typically accompanied by cloud cover). The Copernicus programme does not produce the intelligence; it produces the data that the intelligence is produced from.

The United States Geological Survey's Earth Resources Observation and Science (EROS) Center, located in Sioux Falls, South Dakota, acquires, processes, and distributes Landsat imagery. The Landsat programme, operational since 1972, is the longest continuous satellite Earth observation record. Landsat 8 (launched 2013) and Landsat 9 (launched 2021) provide 30-metre multispectral and 15-metre panchromatic imagery with a 16-day revisit cycle. The EROS Center processes the raw data into analysis-ready products (Level-1, Level-2 surface reflectance) and distributes them freely through EarthExplorer and the LandsatLook viewer. The intelligence produced from the imagery is produced by the Survey's hazard, resource, and environmental programs, and by external users. The EROS Center does not produce the intelligence; it produces the data.

The distinction between remote sensing and geospatial intelligence is the distinction between data and assessment. A remote sensing specialist processes imagery; a geospatial intelligence analyst produces assessments. The remote sensing specialist's product is analysis-ready imagery; the analyst's product is an assessment that informs a decision. The two functions are related — the analyst depends on the specialist's products — but they are distinct, and confusing them undermines both.

The confusion arises most frequently in offices that define themselves by their remote sensing capability rather than by their analytical purpose. An office that defines itself as "the satellite imagery processing centre" will be measured by its data throughput — the number of scenes processed, the volume of data distributed, the uptime of the processing system. An office that defines itself as "the geospatial intelligence office" will be measured by its analytical output — the quality of its flood assessments, the accuracy of its crop forecasts, the timeliness of its infrastructure alerts. The first office produces data; the second produces intelligence. The government that establishes a geospatial intelligence office should be clear about which office it is establishing, because the two offices have different complements, different budgets, different success measures, and different relationships with their consumers.

The geospatial office's value to the government is its analytical output, not its data throughput. An office that processes a million satellite scenes a year but produces no intelligence is a data-processing facility. An office that processes a thousand scenes a year but produces intelligence that informs flood response, granary procurement, infrastructure protection, and mineral exploration is an

intelligence office. The first office's budget is justified by its data distribution; the second office's budget is justified by the value of its intelligence to its consumers. The office that maintains the distinction between remote sensing and geospatial intelligence produces intelligence; the office that blurs it produces data and calls it intelligence.

3.7 The Scope Creep Pattern

The boundaries between geospatial intelligence and its adjacent disciplines are eroded by a pattern that recurs across government geospatial agencies: scope creep. Scope creep is the gradual expansion of the office's work beyond its analytical mandate to include the functions of its adjacent disciplines. The pattern begins with a consumer request that falls on the boundary: a request for a map of existing conditions (cartography), a request for a parcel survey (surveying), a request for a population projection (statistics), a request for a planning recommendation (spatial planning), or a request for raw imagery (remote sensing). The office accepts the request, reasoning that the work is spatial and that the office has the capability. The acceptance establishes a precedent, and the precedent becomes an expectation. The consumer returns with a similar request, and the office accepts again. Over time, the office's workload shifts from analytical work to the adjacent-discipline work, and the analytical output declines.

The scope creep pattern has been observed in multiple government geospatial agencies. A mapping agency that was established to produce topographic maps gradually expands into hazard assessment, because its consumers request hazard maps. A geological survey that was established to produce geological maps gradually expands into environmental monitoring, because its consumers request environmental data. An intelligence office that was established to produce assessments gradually expands into data processing, because its consumers request raw data. In each case, the office's analytical output declines as the adjacent-discipline work expands, and the office's consumers receive less intelligence and more data.

The pattern is difficult to resist because each individual request is reasonable. A map of existing conditions is a spatial product, and the office has the capability to produce it. A parcel survey is a spatial task, and the office has the equipment. A population projection is a statistical task, and the office has the data. The individual requests are not the problem; the pattern is the problem. The office that accepts the individual requests without recognizing the pattern will find its analytical output declining, and the decline will be difficult to reverse because the consumers have come to expect the adjacent-discipline work.

The intake function, treated in the next section, is the mechanism that prevents scope creep. The intake function applies a consistent test to every request and declines the requests that fall outside the office's analytical mandate. The test is not complicated, but it requires the discipline to apply it consistently, and the discipline is maintained by the office manager, who has the authority to decline requests and the responsibility to maintain the office's analytical focus.

3.8 Maintaining the Boundaries in Practice

The discipline boundaries are maintained in practice through the office's intake function. The intake function determines which requests the office accepts and which it declines. The function is

maintained by the office manager, who has the authority to decline requests, and it is documented in the office's procedures so that the criteria for acceptance and decline are transparent.

The intake function applies a test: is the request for an analytical assessment, or is it for a non-analytical spatial product? The test distinguishes between requests that require the office's analytical capability (modeling, validation, assessment) and requests that require only the office's spatial data capability (mapping, surveying, data distribution). An analytical request is accepted; a non-analytical request is declined and referred to the appropriate institution.

A request for a flood exposure assessment — where the office runs a hydraulic model, overlays the inundation with the cadastre, and produces an exposure report — is an intelligence request. The request requires the office's analytical capability: the hydrology plug-in, the hydraulic model, the validation, and the exposure assessment. The request is accepted.

A request for a map of existing conditions, where no analysis is required, is a cartographic request. The request requires the office's cartographic capability but not its analytical capability. The request is declined and referred to the mapping agency.

A request for a survey of a specific parcel is a surveying request. The request requires licensed surveying work, which the office is not authorized to perform. The request is declined and referred to the cadastral authority.

A request for population projections is a statistics request. The request requires demographic modeling, which is the statistics agency's function. The request is declined and referred to the statistics agency.

A request for raw satellite imagery is a remote sensing data request. The request requires data distribution, not analysis. The request is declined and referred to the data provider (the Copernicus programme, the EROS Center, or the commercial imagery provider).

A request for a land-use recommendation is a planning request. The request requires planning judgment, which is the planning authority's function. The request is declined and referred to the planning authority. The office may provide the intelligence that informs the recommendation, but it does not make the recommendation.

The intake function is the office's first line of defense against scope creep. It is maintained by the office manager, who applies the test consistently and documents the decisions. The documentation ensures that the criteria for acceptance and decline are transparent and that the office's consumers understand what the office will and will not do. The documentation also provides the audit trail that the national audit authority can examine to verify that the office is maintaining its analytical focus.

The intake function is not a one-time decision. The office's consumers will continue to submit requests that fall on the boundary, and the office manager will continue to apply the test. The consistency of the application is the discipline, and the discipline is the mechanism that maintains the boundaries over time. An office that applies the test inconsistently — accepting boundary requests from some consumers and declining them from others — will erode the boundaries and will create the perception that the office's mandate is negotiable. An office that applies the test consistently will maintain the boundaries and will establish the expectation that the office's mandate is fixed.

3.9 Worked Example: Applying the Intake Test

The office receives five requests in a single morning. The first is from the emergency management authority, requesting a flood inundation assessment for the Riverbourne Basin following 180 millimetres of rainfall over 48 hours. The request specifies the area of interest, the required products (inundation map, exposure report, evacuation guidance), and the delivery timeline (within 4 hours). The office manager applies the intake test: the request requires the office's analytical capability (the hydrology plug-in, the hydraulic model, the validation, the exposure assessment). The request is accepted. The office manager activates the incident response procedure (section 13.6) and assigns the hazard specialist to produce the assessment.

The second request is from the strategic reserve authority, requesting the weekly crop condition bulletin for the current growing season. The request is a standing request, received every Monday morning, and the office produces the bulletin every Wednesday by close of business. The office manager applies the intake test: the request requires the office's analytical capability (the crop model plug-in, the vegetation indices, the yield forecast, the parcel stratification). The request is accepted as a standing request.

The third request is from the energy sector entity, requesting a map of the transmission corridor between Substation A and Substation B, showing the current vegetation cover. The request specifies the corridor (a 50-kilometre transmission line), the required product (a map), and the format (GeoTIFF). The office manager applies the intake test: the request is for a map of existing conditions, where no analysis is required. The request does not require the office's analytical capability. The request is declined and referred to the national mapping agency, which maintains the vegetation cover data and can produce the map.

The energy sector entity responds with a modified request: not just a map of existing conditions, but an assessment of whether the vegetation cover constitutes an encroachment risk to the transmission corridor, with a vulnerability score for each segment. The office manager applies the intake test to the modified request: the request now requires the office's analytical capability (the change detection plug-in, the encroachment detection algorithm, the vulnerability scoring). The modified request is accepted and assigned to the hazard specialist.

The fourth request is from the ministry of agriculture, requesting a survey of a specific 50-hectare parcel to verify the boundary for a land consolidation proceeding. The office manager applies the intake test: the request is for a cadastral survey, which requires licensed surveying work. The office is not authorized to perform cadastral surveys. The request is declined and referred to the cadastral authority.

The fifth request is from a commercial minerals exploration company, requesting raw Sentinel-2 satellite imagery for a specific area of interest. The office manager applies the intake test: the request is for raw satellite imagery, which is a data distribution request, not an analytical request. The request is declined and referred to the Copernicus Open Access Hub, where Sentinel-2 data is available free of charge. The office manager notes that if the company requests a prospectivity assessment for the area, that request would be accepted as a commercial commission under the commercial release procedure (section 13.4).

The five requests illustrate the intake test's operation. Two requests (the flood assessment and the crop bulletin) are accepted because they require the office's analytical capability. One request (the

encroachment assessment) is accepted after modification because the modified request requires the office's analytical capability. Two requests (the parcel survey and the raw imagery) are declined and referred to the appropriate institutions because they do not require the office's analytical capability. The intake function has maintained the office's analytical focus by applying a consistent test to every request, and the documentation of the decisions provides the audit trail that verifies the office's discipline.

Chapter 4 — The Geospatial Officer as a Professional Category

4.1 The Role Family

The geospatial office described in this Handbook has six seats: the office manager, the technical lead, the data pipeline operator, the production operator, the hazard specialist, and the resource economics specialist. These six seats form a single role family because they share a common technical foundation — the spatial register, the pipeline, the classification framework — and because the work of each seat depends on the work of the others. The role family is distinct from the generalist administrative stream, because the work requires specialized technical knowledge that the generalist stream does not provide. It is also distinct from the IT professional stream, because the work requires analytical and cartographic judgment that the IT stream does not provide.

The six-seat structure is the complement appropriate to a small geospatial office. National geospatial agencies have complements ranging from a dozen to several thousand, organized in role families that reflect the agency's mission and the jurisdiction's civil service traditions. The United States Geological Survey employs approximately 8,000 staff across its hazard, resource, environmental, and mapping programs. The Survey's scientific staff are organized in occupational series: the 0400 series (natural resources management and biological sciences), the 1300 series (physical sciences), the 1350 series (geology), the 1400 series (physics), and the 1700 series (mathematics and statistics). Technical staff are organized in the 1300 and 1500 series, and cartographic staff in the 1370 series (cartography). The Survey's senior leadership is drawn from the Senior Executive Service, which is a government-wide corps of senior executives separate from the occupational series.

Geoscience Australia employs approximately 500 staff, organized in the agency's geoscience series, spatial information series, and corporate services series. The geoscience series covers the agency's geologists, geophysicists, and environmental scientists; the spatial information series covers the agency's remote sensing specialists, GIS analysts, and cartographers; the corporate services series covers the agency's IT, finance, HR, and communications staff. The agency's senior leadership is drawn from the Senior Executive Service of the Australian Public Service.

The French Institut national de l'information géographique et forestière employs approximately 1,800 staff, organized under the French civil service framework. The agency's technical staff belong to the corps of ingénieurs des travaux géographiques et cartographiques (engineers of geographic and cartographic works) and the corps de techniciens géographiques (technician corps). The agency's scientific staff belong to the corps of ingénieurs de l'agriculture et de l'environnement (engineers of agriculture and environment). The agency's administrative staff belong to the corps of administrateurs civils (civil administrators). The French corps system is a career-long structure: an officer enters a corps through a competitive examination, progresses through the corps's ranks, and may move between corps through internal mobility procedures.

The six-seat structure treated in this Handbook is the complement appropriate to a small geospatial office. It scales by adding seats to the family and by adding layers to the hierarchy, not by replacing the family with a different structure. An office that grows from six to twelve seats adds a second specialist in each domain track, a second data pipeline operator, and a second production operator, and the office manager and technical lead remain. An office that grows from twelve to

twenty-four seats adds a deputy office manager and a deputy technical lead, and the structure becomes a two-layer hierarchy with the office manager and technical lead at the top and the deputies and operators at the bottom. The role family structure accommodates the growth without restructuring.

4.2 The Rank Structure

The geospatial role family follows the general pattern of government rank structures, with an operational tier and an executive tier. The operational tier comprises the data pipeline operator, the production operator, the two specialists, and the technical lead. The executive tier comprises the office manager. The operational tier conducts the technical work; the executive tier manages the work and bears accountability for it. This division ensures that the manager is not so absorbed in operational work that the management function is neglected.

The two-tier structure is standard across mature civil services. The United States federal government separates the General Schedule (GS), which covers most professional positions in 15 grades (GS-1 through GS-15), from the Senior Executive Service (SES), which covers senior management positions. Geospatial professionals in the United States Geological Survey are typically classified in the GS-0400, GS-1300, or GS-1700 series, with entry-level positions at GS-7 to GS-9, mid-career positions at GS-11 to GS-13, and senior positions at GS-14 to GS-15. The Survey's center directors and senior scientists are typically in the SES. The Canadian federal government separates the occupational groups (such as the Research Scientist group, the Biologist group, and the Computer Systems group) from the Executive group (EX-1 through EX-5). The French civil service separates the corps of engineers and technicians (categories A and B) from the corps of administrators (category A+). The Singapore civil service separates the management associates (who are groomed for senior leadership) from the generalist officers (who progress through the operational ranks).

Within the operational tier, the ranks are differentiated by experience and responsibility. The data pipeline operator and the production operator are the entry-level operational ranks. They require technical competence — the ability to operate the acquisition desk, the processing engine, the product studio, and the catalog and archive submodules — but they do not require the analytical depth of the specialist ranks. The operators are the officers who keep the pipeline running: they ingest the data, configure the processing runs, produce the maps, and archive the products. Their work is the operational backbone of the office, and the quality of their work determines the quality of the pipeline's output.

The two specialists are the senior operational ranks. They require domain expertise in their respective tracks — the hazard specialist in hydrology and infrastructure monitoring, the resource economics specialist in agriculture and mineral prospectivity — and the analytical judgment to produce intelligence that is credible to the office's consumers. The specialists are the officers who produce the office's intelligence: they run the models, validate the results, interpret the outputs, and produce the assessments. Their work is the analytical core of the office, and the quality of their work determines the quality of the office's intelligence.

The technical lead is the bridge between the operational and executive tiers. The technical lead holds an operational rank but exercises a bridging function that includes pipeline configuration, model validation, and specialist support. The technical lead is the office's senior technical

professional, and the lead's role is to ensure that the office's technical work meets the office's technical standards. The technical lead is the officer whom the office manager relies on for technical advice, and the officer whom the specialists rely on for technical support. The technical lead is also the officer who configures the modeling workbench, sets the processing parameters, and validates the model runs — the technical decisions that determine the quality of the office's output.

The office manager is the executive tier. The manager is accountable for the office's production, its compliance with standards, its financial administration, and its maintenance of the commercial firewall. The manager is the sole authority for commercial releases, and the authority cannot be escalated or delegated. The manager reports to the desk chief and is evaluated on the office's performance against the standards set out in this Handbook.

The detailed rank structure, with grade levels and progression criteria, is treated in Chapter 30, which constitutes the formal submission to the qualifications authority for the establishment of the geospatial role family.

4.3 Career Paths

Three principal career paths run through the role family, each reflecting a different trajectory. The pipeline path enters as a data pipeline operator, progresses to the technical lead, and may advance to the office manager. Officers on this path develop deep technical knowledge of the pipeline and the systems, and they become the technical leads who configure the modeling workbench, the processing parameters, and the validation protocols. The pipeline path is the path for officers who are drawn to the technical infrastructure of the geospatial office: the data acquisition, the processing chains, the base layer maintenance, the system configuration. A data pipeline operator who progresses on the pipeline path acquires additional certifications in the processing engine and the spatial data infrastructure, takes on increasing responsibility for the pipeline's configuration, and eventually assumes the technical lead role with primary responsibility for the pipeline's technical integrity.

The production path enters as a production operator, progresses to a specialist track through additional training and certification, and may advance to the technical lead or the office manager. Officers on this path develop deep knowledge of the production and cartographic function, and they become the specialists who understand how analytical outputs are transformed into delivered products. The production path is the path for officers who are drawn to the communication dimension of the geospatial office: the cartographic design, the classification marking, the format conversion, the archival. A production operator who progresses on the production path acquires additional training in a domain (hydrology, agriculture, geology, or infrastructure monitoring), transitions to the specialist rank in that domain, and eventually becomes a specialist who produces the domain's intelligence while understanding the production constraints that shape the product's form.

The domain path enters through a domain-specific recruitment, progresses through the specialist rank, and may advance to the technical lead or the office manager. Officers on this path develop deep domain expertise, and they become the specialists who produce the flood models, the crop assessments, the infrastructure monitoring, and the mineral prospectivity studies. The domain path is the path for officers who enter the geospatial profession through a domain science rather than through the geospatial technology: the hydrologist who learns GIS, the geologist who learns remote sensing, the agronomist who learns crop modeling. A domain-path officer enters as a specialist, acquires the

geospatial certifications, and progresses through the specialist rank to the technical lead or the office manager.

The three paths converge at the technical lead and the office manager ranks. This convergence is deliberate. Officers who reach the senior ranks should have both depth and breadth, and the convergence ensures that the technical lead and the office manager have been exposed to multiple perspectives on the office's work. An officer who has only worked on the pipeline may not understand the specialists' analytical challenges. An officer who has only worked as a specialist may not understand the pipeline's operational constraints. The convergence at the senior ranks produces officers who understand both, and the understanding is the condition of effective senior leadership in a geospatial office where the pipeline and the analytical work are interdependent.

The multi-path model is standard in national geospatial agencies. The United States Geological Survey's career paths include the 1310 series (geology), the 1350 series (geology), the 1370 series (cartography), the 1315 series (hydrology), and the 0401 series (natural resources management), each with its own progression. The Survey encourages movement between series as a means of developing breadth, and the Survey's senior scientists have typically worked in multiple series over their careers. Geoscience Australia's career paths include the geoscience series, the spatial information series, and the corporate services series, with progression within each and movement between them. The multi-path model produces officers with both depth and breadth, and it is the model that the geospatial role family adopts.

4.4 Certification and Qualifications

The qualifications framework for the geospatial role family has three layers: academic qualifications, technical certifications, and professional certifications.

The academic qualifications are the university degrees or equivalent technical qualifications that establish the officer's foundational knowledge. The operators require a technical qualification in geospatial technology, geographic information science, or a related field — typically a two-year technical diploma or a four-year bachelor's degree. The specialists require a university degree in their domain: the hazard specialist requires a degree in hydrology, civil engineering, environmental science, or a related field; the resource economics specialist requires a degree in agriculture, geology, earth science, or a related field. The technical lead requires a university degree in geospatial science, geographic information science, or a related field, with additional technical certifications. The office manager requires a university degree and the management qualifications specified by the government's executive rank requirements, which may include a master's degree in public administration, business administration, or a related field.

The technical certifications establish the officer's competence in the specific tools and methods the office uses. The operators require certifications for the submodules they operate: the data pipeline operator requires certifications for the acquisition desk and the processing engine; the production operator requires certifications for the product studio and the catalog and archive. The specialists require certifications for the modeling workbench and their domain plug-ins: the hazard specialist requires certifications for the hydrology plug-in and the change detection plug-in; the resource economics specialist requires certifications for the crop model plug-in and the prospectivity plug-in. The technical lead requires certifications for all six submodules, with additional certifications in

system configuration, validation methodology, and quality assurance. The office manager requires certifications for workflow management, classification management, and the commercial release authorization procedure.

The technical certifications are awarded by the administrative training authority (treated in Chapter 31) on completion of the relevant course and the passing of the assessment. The certifications are maintained through continuing education: the officer must complete refresher courses, update courses, or advanced courses on a periodic basis to maintain the certification. An officer who does not maintain the certification loses the qualification to operate the submodule or the domain plug-in, and the loss of the qualification means the officer must be reassigned or recertified before resuming the work.

The professional certifications are offered by national and international professional bodies. The American Society for Photogrammetry and Remote Sensing (ASPRS), founded in 1934, offers the Certified Photogrammetrist (CP) and Certified Mapping Scientist (CMS) certifications. The GIS Certification Institute (GISCI), established in 2004, offers the GIS Professional (GISP) certification, which requires a combination of education, experience, and contributions to the profession. The Surveying and Spatial Sciences Institute (SSSI), the professional body for Australia and New Zealand, offers the Certified Professional in Geographic Information Systems (CPGIS) certification. The International Society for Photogrammetry and Remote Sensing (ISPRS), founded in 1910, provides the international scientific framework for the discipline through its commissions and working groups. The Royal Institution of Chartered Surveyors (RICS), established in 1868, offers the Chartered Surveyor qualification with a geospatial pathway.

Professional certifications are encouraged but not required. They establish the officer's standing in the broader profession and provide a community of practice that sustains the professional identity over a career. The United States Geological Survey encourages its scientists to hold professional certifications through ASPRS and GISCI. Geoscience Australia encourages professional certification through SSSI. The Ordnance Survey encourages professional certification through RICS. In each case, the agency supports the certification by providing study leave, reimbursing examination fees, and recognizing the certification in the agency's performance management system.

The certification framework is designed to be compatible with the government's general civil service framework, so that geospatial officers can move between the geospatial role family and the generalist administrative stream or the IT professional stream. This compatibility matters because it ensures that the geospatial role family does not become a closed career. Officers who enter the role family can exit it if the office is dissolved or if their career aspirations change, and the office can recruit from the generalist and IT streams when it needs officers with different backgrounds.

4.5 Language Requirements

The language requirements of the geospatial role family arise from two demands. The government's operating language is the language in which the government conducts its business, and the geospatial office must produce its products in that language. The international character of the geospatial profession means that the scientific literature, the technical documentation, and the professional discourse are conducted in a small number of international languages — predominantly

English — and the geospatial officer must be able to read and understand them. The language requirements are the intersection of these two demands, and they vary by rank and by role.

The operators require reading and writing proficiency in the government's operating language. The operators must read the system interfaces (which are typically in English, regardless of the government's operating language, because most geospatial software is developed in English), write the records (which are in the government's operating language), and communicate with their colleagues (in the operating language). The operators do not require speaking proficiency in an international language, because the operators do not brief external audiences or read the international scientific literature.

The specialists require reading and writing proficiency in the operating language, plus reading proficiency in at least one international language (typically English). The specialists must read the scientific literature (which is predominantly in English: the ISPRS International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, the International Journal of Applied Earth Observation and Geoinformation, the Journal of Hydrology, the journal Remote Sensing of Environment), the technical documentation (which is in English for most geospatial software), and the professional discourse (which is conducted in English at international conferences and in professional bodies). The specialists must also write the intelligence products (which are in the operating language) and communicate with their consumers (in the operating language).

The technical lead requires reading and writing proficiency in the operating language, plus reading and writing proficiency in at least one international language. The technical lead must read and write the technical documentation (which is in English), participate in the professional discourse (which is in English), and read the international scientific literature. The technical lead may also need to communicate with international counterparts — the technical leads of other national geospatial agencies, the technical committees of international standards bodies, the technical working groups of international Earth observation programmes — and this communication is conducted in English.

The office manager requires reading, writing, and speaking proficiency in the operating language, plus reading proficiency in at least one international language. The manager must brief the desk chief and the bureau chief (in the operating language), present the office's performance at the monthly and quarterly reviews (in the operating language), and read the international literature (in English). The manager may also need to represent the office at international forums — the United Nations Committee of Experts on Global Geospatial Information Management, the Group on Earth Observations, the International Organization for Standardization's Technical Committee 211 on Geographic Information — and this representation requires speaking proficiency in English.

The comparative practice confirms the dual language demand. The United States Geological Survey operates in English and requires its scientists to read the international scientific literature, which is predominantly in English. The French Institut national de l'information géographique et forestière operates in French and requires its engineers to read the international literature in English; the agency's scientific publications are produced in French, but the agency's participation in international forums (the EuroGeographics association, the Permanent Committee on GIS Infrastructure for Asia and the Pacific) is conducted in English. The Singapore civil service operates in English and requires its officers to be proficient in at least one other official language (Malay, Mandarin, or Tamil), reflecting Singapore's multilingual society.

4.6 Payscale Considerations

Geospatial skills are in demand in the private sector, in the resource industries, in the technology sector, and in the consulting industry. The demand means that the government's standard payscale may not be sufficient to attract and retain geospatial officers. The payscale modifier is the mechanism by which the government adjusts the payscale for specific roles where labour-market conditions require it. The geospatial role family is a candidate for a payscale modifier in most jurisdictions.

The labour-market demand for geospatial skills has several drivers. The resource industries — mining, oil and gas — employ geospatial professionals for exploration support, resource estimation, and environmental monitoring. The technology sector — Google, Microsoft, Amazon, Esri — employs geospatial professionals for mapping platforms, location-based services, and Earth observation data processing. The consulting industry — firms such as Deloitte, Accenture, and Booz Allen Hamilton — employs geospatial professionals for government contracting. The commercial Earth observation sector — Maxar, Planet, Airbus — employs geospatial professionals for imagery analysis and product development. These sectors offer salaries that may be substantially higher than the government's standard payscale, and the salary differential makes it difficult for the government to attract and retain geospatial officers.

The United States federal government uses special pay rates for certain scientific and technical positions where the labour market is competitive. The special pay rate system, authorized by 5 U.S.C. § 5305, allows the Office of Personnel Management to establish higher pay rates for specific occupations and localities. Geospatial positions in the United States Geological Survey and the National Geospatial-Intelligence Agency have received special pay rate adjustments in the past, particularly for positions requiring advanced degrees or specialized certifications. The Australian government uses market-related pay for certain technical positions under the agency's enterprise agreements, and Geoscience Australia has used market-related pay to attract and retain geospatial professionals in competition with the resource industries. The Singapore civil service uses market-related pay across the service, benchmarked to private-sector comparators, and geospatial positions in the Singapore Land Authority and other agencies are compensated at market rates.

The payscale modifier is applied as a percentage adjustment to the base payscale, with the percentage calibrated to the labour-market conditions in the jurisdiction. The modifier is not a permanent feature of the role family; it is reviewed periodically — typically every two to three years — and adjusted or removed as labour-market conditions change. A payscale modifier that is no longer needed is a fiscal cost that the government should not bear. A payscale modifier that is too low is a recruitment and retention problem that the office should not have to manage. The review is conducted by the office manager, who assesses the labour-market conditions (turnover rates, time-to-fill, competitor salaries) and recommends the modifier's adjustment, and the qualifications authority, which ratifies the modifier.

The payscale modifier should be applied with care. A modifier that is applied too broadly — to all geospatial positions regardless of the labour-market conditions for the specific role — creates fiscal costs that are not justified by the labour-market evidence. A modifier that is applied too narrowly — only to the specialist ranks, not to the operators — may fail to address the recruitment and retention challenges at the operational level, where the labour-market competition is often strongest (because the operators' skills are more transferable to the private sector than the specialists'

domain expertise). The modifier should be applied to the roles where the labour-market evidence supports it, and the evidence should be documented so that the qualifications authority and the national audit authority can verify the modifier's justification.

4.7 The Professional Identity

The professional identity of the geospatial officer is the identity the officer carries as a member of the geospatial profession. It is not a matter of job titles or rank insignia; it is a matter of the officer's self-understanding as a practitioner of a specific discipline, with the responsibilities and standards that the discipline imposes. An officer who identifies as a geospatial intelligence professional will produce intelligence. An officer who identifies as a spatial professional without a disciplinary anchor will produce whatever the reporting relationship pressures, and the product will not be intelligence.

The professional identity matters because it is the individual expression of the discipline boundaries treated in Chapter 3. The discipline boundaries — the distinction from surveying, cartography, spatial planning, statistics, and remote sensing — are maintained through institutional arrangements, but the institutional arrangements are only as effective as the officers who operate within them. An officer who understands the boundaries and identifies with the analytical function will maintain the boundaries through the intake function and the daily work. An officer who does not understand the boundaries, or who identifies with a different function, will erode the boundaries through the gradual acceptance of adjacent-discipline work. The professional identity is the internal mechanism that maintains the discipline boundaries when the external mechanisms (institutional placement, intake procedures) are not sufficient.

The professional identity is established through three mechanisms. The training pipeline (treated in Chapter 31) immerses the officer in the discipline's methods, standards, and literature through the foundational course, the submodule courses, and the domain courses. The qualifications framework (treated in section 4.4) certifies the officer's competence in the discipline through the academic qualifications, the technical certifications, and the professional certifications. The professional bodies — the American Society for Photogrammetry and Remote Sensing (founded 1934), the GIS Certification Institute (established 2004), the Surveying and Spatial Sciences Institute (the professional body for Australia and New Zealand), the International Society for Photogrammetry and Remote Sensing (founded 1910), and the Royal Institution of Chartered Surveyors (established 1868) — provide the officer with a community of practice that sustains the identity over a career.

The professional bodies contribute to the professional identity in several ways. They publish peer-reviewed journals (PE&RS — Photogrammetric Engineering & Remote Sensing for ASPRS; the International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences for ISPRS; Geomatics Professional for SSSI) that disseminate the discipline's methods and findings. They organize annual conferences (the ASPRS Annual Conference, the ISPRS Congress, the SSSI Locate conference) where the officers present their work, learn from their peers, and build the professional relationships that sustain the community. They develop and maintain the professional certifications (ASPRS Certified Photogrammetrist, GISCI GIS Professional, SSSI Certified Professional) that establish the officer's standing in the profession. They participate in the international standards bodies (ISO/TC 211, the Open Geospatial Consortium) that develop the technical standards the officer uses.

National geospatial agencies that invest in their officers' professional identity — through training, through certification, through professional body membership — produce officers who maintain the discipline boundaries discussed in Chapter 3 and who produce intelligence that is credible to the office's consumers. The United States Geological Survey supports its scientists' participation in ASPRS and ISPRS through membership fee reimbursement, conference travel funding, and time-off for professional service. Geoscience Australia supports its scientists' participation in SSSI and ISPRS through similar mechanisms. The Ordnance Survey supports its staff's participation in RICS. Agencies that do not invest in the professional identity produce officers who drift toward whatever function the reporting relationship pressures, and their intelligence output declines.

4.8 Worked Example: Recruiting and Developing a Hazard Specialist

The office has a vacancy for a hazard specialist following the promotion of the previous incumbent to technical lead. The office manager initiates the recruitment process.

The position description specifies the following qualifications. Academic: a university degree in hydrology, civil engineering, environmental science, or a related field. A master's degree is preferred but not required. Technical: certifications for the modeling workbench, the hydrology plug-in, and the change detection plug-in, or the willingness to acquire these certifications within 12 months of appointment. Professional: professional certification (ASPRS Certified Mapping Scientist, GISCI GIS Professional, or equivalent) is encouraged but not required. Language: reading and writing proficiency in the government's operating language, plus reading proficiency in English for the international scientific literature. Experience: a minimum of three years of experience in hydrological modeling, flood risk assessment, or infrastructure monitoring, in a government agency, a consulting firm, or a research institution.

The recruitment attracts 14 applications. The office manager and the technical lead review the applications and shortlist four candidates. The shortlist includes a hydrologist from the national meteorological service with five years of flood forecasting experience, a civil engineer from a consulting firm with four years of flood risk assessment experience, a recent master's graduate in environmental science with internship experience at the geological survey, and a GIS analyst from the municipal planning office with six years of experience in hazard mapping but no formal hydrology training.

The office manager and the technical lead interview the four candidates. The interviews cover the candidate's technical knowledge (hydrological modeling, HEC-RAS, MIKE 21, TUFLOW), the candidate's analytical judgment (how would you validate a flood model against observed events?), the candidate's understanding of the discipline boundaries (what would you do if a consumer requested a map of existing conditions rather than a flood assessment?), and the candidate's professional identity (what professional bodies do you belong to? what was the last conference you attended?).

The hydrologist from the national meteorological service is the strongest candidate. The hydrologist has five years of flood forecasting experience using the HEC-RAS model, has published in the Journal of Hydrology, is a member of the International Association of Hydrological Sciences, and has a clear understanding of the distinction between producing intelligence and conducting

operations. The hydrologist does not have the modeling workbench certifications, but the office manager determines that the certifications can be acquired within the 12-month period through the administrative training authority's courses. The hydrologist is offered the position and accepts.

The new hazard specialist's development plan covers the first 12 months. Months 1-2: complete the foundational course at the administrative training authority, covering the pipeline architecture, the information system, the classification framework, and the office's procedures. Months 3-4: complete the modeling workbench course and the hydrology plug-in course, and pass the certification assessments. Months 5-6: complete the change detection plug-in course and pass the certification assessment. The specialist begins producing flood model runs under the technical lead's supervision during this period. Months 7-8: begin independent model runs, with the technical lead validating a sample of the runs. The specialist begins producing energy infrastructure monitoring runs under the technical lead's supervision. Months 9-10: complete independent model runs in both domains, with the technical lead validating a sample. The specialist begins producing the weekly crop condition bulletin as cross-training (the resource economics specialist is on leave). Months 11-12: the specialist is fully operational in both domains, producing flood model runs and energy infrastructure monitoring runs independently, with the technical lead's sampling validation. The specialist applies for the ASPRS Certified Mapping Scientist certification, with the office manager's support.

The development plan illustrates the role family's structure in practice. The new specialist enters through the domain path (hydrology), acquires the technical certifications (modeling workbench, hydrology plug-in, change detection plug-in) through the training pipeline, develops the professional identity through the professional body membership (ASPRS, IAHS), and progresses from supervised to independent work over 12 months. The development plan is the operational expression of the role family's career paths, certification framework, and professional identity, and it is the mechanism by which the office develops the officer from a new hire to a fully operational specialist.

Chapter 5 — The Commercial Firewall as Foundational Principle

5.1 The Principle

The commercial firewall is the institutional separation between a geospatial office's commercial arm and its public-service arm. The commercial arm sells mineral intelligence to external customers and generates the revenue that sustains the public-service arm. The public-service arm produces flood risk, food security, and energy infrastructure intelligence for government consumers. The firewall ensures that the commercial arm's customer orientation does not distort the public-service arm's intelligence, and that the public-service arm's public-interest orientation does not distort the commercial arm's pricing and product decisions.

The commercial arm is a legitimate and substantial function of the office. It produces mineral intelligence that is valued by its customers — mining and exploration companies that use prospectivity maps to reduce their exploration risk — and that generates the revenue that funds the flood risk modelling, the crop condition assessments, and the infrastructure monitoring that the public-service arm produces for the government. The firewall does not subordinate the commercial arm to the public-service arm. It separates them, so that each can operate according to its own logic without distorting the other.

The need for the firewall arises from a structural conflict of interest that is inherent in the revenue-neutral model. The commercial arm's revenue depends on the goodwill of its customers — the mining and exploration companies. The public-service arm's intelligence may affect the interests of those same companies — a flood risk assessment may constrain development in an area of mineral interest, an infrastructure monitoring report may identify risks at a facility owned by a company that also purchases mineral intelligence. If the two arms are not separated, the commercial arm's financial interest in keeping its customers content creates pressure on the public-service arm to soften, delay, or withhold intelligence that might displease those customers. The firewall is the mechanism that prevents this pressure from operating.

The firewall is maintained through four mechanisms. The financial mechanism uses separate charge codes, separate accounts, and separate reporting, so that the commercial arm's revenue and the public-service arm's expenditure are visible and auditable as distinct streams. The classification mechanism requires every commercial product to carry a classification marking that restricts it to the customer and the office, and every public-service product to carry a classification marking that reflects its sensitivity, so that the two arms' products are not commingled in the archive or in the distribution system. The authority mechanism reserves commercial release authorization to the office manager alone, and the authority cannot be escalated or delegated. The audit mechanism requires the national audit authority to examine the commercial arm's accounts and the office's compliance with the firewall, so that the firewall's maintenance is verified by an external authority.

5.2 Comparative Models

The revenue-neutral model with a commercial firewall is exemplified by national geological surveys that fund public geoscience through mineral-data sales. The United States Geological Survey

sells mineral resource data and custom products through its Earth Resources Observation and Science (EROS) Center, located in Sioux Falls, South Dakota. The EROS Center was established in 1973 as the federal government's centre for land remote sensing data acquisition, processing, and distribution. It now holds the world's largest collection of civilian remotely sensed data, including the Landsat archive dating to 1972. The revenue from custom product sales and value-added services supports the Survey's public geoscience programs. The firewall is maintained through the Survey's financial architecture, which separates commercial revenue from appropriated funding, and through the Survey's publication policy, which distinguishes public data (freely available) from commercial products (sold on a cost-recovery basis).

Geoscience Australia sells mineral intelligence through its minerals division. The division produces prospectivity maps, resource estimates, and custom geological studies for mining and exploration companies operating in Australia. The revenue supports the agency's public geoscience programs, including the geological mapping of the Australian continent (conducted under the Exploring for the Future programme, launched in 2016), the national seismic hazard assessment, and the groundwater assessment programme. The firewall is maintained through the Geoscience Australia Act 2000, which establishes the conditions under which the agency may sell products and the conditions under which the revenue may be used, and through the agency's internal financial controls, which separate the commercial revenue from the appropriated funding.

The British Geological Survey (BGS), established in 1835 as the Geological Survey of Great Britain, is the world's oldest national geological survey. BGS sells commercial geoscience services through its commercial division, which provides custom geological reports, site investigations, and geotechnical assessments to the construction, mining, and energy sectors. The revenue supports the Survey's public science, including the national geological map of the United Kingdom (the DiGMapGB digital geological map series), the national groundwater monitoring network, and the British National Seismological Database. BGS operates as a research centre of the Natural Environment Research Council (NERC), and its commercial activities are governed by NERC's trading fund regulations.

National mapping agencies with cost-recovery arms provide further examples. The Ordnance Survey of Great Britain, established in 1791, became a trading fund in 1999 under the Trading Funds Act 1973. The trading fund framework gives the Ordnance Survey commercial freedom — it can set prices, negotiate contracts, and invest in its business — while requiring it to maintain the public mapping function that is its statutory obligation. The Ordnance Survey's Public Sector Mapping Agreement (PSMA), established in 2010, provides free access to Ordnance Survey data to all public sector bodies in England and Wales, while commercial users pay commercial rates. The PSMA is the firewall's operational expression: the public sector receives the data for free, the commercial sector pays, and the two sets of terms are separated by the trading fund's financial architecture. The Singapore Land Authority operates a similar model, with commercial data sales funding public mapping functions.

The comparative record also includes negative examples: agencies that blurred the firewall and saw their public-service functions distorted by commercial pressure. The cases are not always public, because the agencies involved have an incentive to conceal the distortion, but the pattern has been documented in audit reports, after-action reviews, and academic studies. The pattern is consistent

across cases. The commercial arm's revenue becomes important to the agency's budget, and the agency's leadership begins to treat the commercial arm's customers as stakeholders whose interests must be considered in the agency's work. The public-service arm's intelligence begins to serve the commercial customers' interests rather than the public interest: a hazard assessment that would constrain a customer's development is delayed; an infrastructure monitoring report that would identify a customer's facility is softened; a mineral prospectivity study that would benefit a customer's competitor is withheld. The agencies that maintain the firewall produce both commercial products and public-service intelligence that are credible in their respective markets. The agencies that erode the firewall produce products that are contested in both markets — the commercial customers distrust the intelligence because they suspect it is shaped by the public-service mandate, and the public-service consumers distrust the intelligence because they suspect it is shaped by the commercial interest.

5.3 The Statutory Basis

The firewall must be maintained as a statutory or quasi-statutory arrangement, not as a managerial discretion. A manager can be reassigned, pressured, or replaced; a statute cannot. A firewall maintained by managerial discretion will be eroded by the first manager who finds it inconvenient, and the erosion will be gradual, invisible, and irreversible. A firewall maintained by statute will survive the managers who find it inconvenient, because the statute cannot be changed without a legislative act, and the legislative act requires public deliberation.

The United States Geological Survey's commercial activities are governed by the Survey's statutory authority, which derives from the Organic Act of 1879 (20 U.S.C. §§ 31-44) and subsequent legislation. The Survey's authority to sell products and services is established by 43 U.S.C. § 1460, which authorizes the sale of maps, publications, and other data products, and by the Federal Technology Transfer Act of 1986 (15 U.S.C. §§ 3701-3714), which authorizes cooperative research and development agreements (CRADAs) with private sector partners. The revenue from these sales is deposited in the Survey's revolving fund, authorized by 43 U.S.C. § 364b, which allows the Survey to retain and use the revenue for its programmes. The statutory framework establishes the conditions under which the Survey may sell, the conditions under which the revenue may be used, and the reporting requirements that make the commercial activities transparent.

The Ordnance Survey's commercial activities are governed by the Ordnance Survey Act 1841 and the Trading Funds Act 1973. The Ordnance Survey Act 1841 established the Ordnance Survey as a government body with the duty to survey and map Great Britain. The Trading Funds Act 1973 allowed the Ordnance Survey to become a trading fund in 1999, which gave it commercial freedom while remaining in government ownership. The trading fund framework requires the Ordnance Survey to cover its costs from a combination of commercial revenue and government contracts, to publish annual accounts, and to be audited by the National Audit Office. The framework also establishes the Public Sector Mapping Agreement, which separates the public sector's free access from the commercial sector's paid access.

Geoscience Australia's commercial activities are governed by the Geoscience Australia Act 2000. The Act establishes the agency as a body corporate with the power to charge for goods and services, to enter into contracts, and to acquire and dispose of property. The Act requires the agency

to prepare an annual report, including financial statements, and to have those statements audited by the Auditor-General. The Act does not establish a trading fund or a revolving fund — the agency's revenue is returned to the Consolidated Revenue Fund, and the agency's budget is appropriated annually — but the agency's internal financial controls separate the commercial revenue from the appropriated funding, and the agency's annual report discloses the commercial revenue separately.

The statutory basis should specify the four mechanisms described in section 5.1. It should require that the commercial arm's revenue is accounted for separately from the public-service arm's expenditure. It should require that commercial products carry classification markings restricting them to the customer and the office. It should reserve commercial release authority to the office manager and prohibit delegation. It should require the national audit authority to examine the commercial arm's accounts and the office's compliance with the firewall. It should specify the conditions under which the firewall may be modified, and the modification should require a legislative act, not a managerial decision.

The statutory basis is not a one-time achievement. It must be maintained through the legislative process, which means that the office's leadership, the office's consumers, and the professional community must advocate for the firewall's continuation when the legislation is reviewed or when amendments are proposed. The firewall's statutory basis has been challenged in several jurisdictions by budgetary pressures (the commercial revenue is attractive to a treasury seeking to reduce appropriated funding), by political pressures (the commercial arm's customers may lobby for reduced public-service scrutiny of their activities), and by institutional reorganizations (a ministry reorganization may merge the geospatial office with another agency that does not have the firewall's culture). The statutory basis survives these challenges only if the office's leadership and the professional community are vigilant in its defence.

5.4 The Financial Architecture

The financial architecture of the firewall comprises separate charge codes, separate accounts, and separate reporting. The commercial arm's revenue is recorded in a commercial revenue account. The public-service arm's expenditure is recorded in a public-service expenditure account. The two accounts are not commingled. The commercial arm's costs — staff time, data acquisition, processing, production — are recorded in a commercial cost account, and the commercial arm's revenue must exceed its costs for the commercial function to be sustainable. The public-service arm's costs are recorded in a public-service cost account, and the public-service arm is funded by the commercial arm's surplus, supplemented by appropriated funding where the commercial surplus is insufficient.

The charge code separation is the operational mechanism that makes the financial separation work in practice. Every task the office performs is assigned a charge code: either a commercial charge code (for work on commercial commissions) or a public-service charge code (for work on public-service products). The data pipeline operator assigns the charge code when the data is acquired (commercial data is acquired under the commercial charge code; public-service data is acquired under the public-service charge code). The specialists assign the charge code when the model runs are configured (commercial model runs use the commercial charge code; public-service model runs use the public-service charge code). The production operator assigns the charge code when the products are produced (commercial products carry the commercial charge code; public-service products carry

the public-service charge code). The charge code is recorded in the catalog and archive submodule, and it is the field that the audit authority uses to verify the financial separation.

The financial architecture includes a smoothing mechanism, because the commercial arm's revenue fluctuates with commodity prices and exploration cycles, while the public-service arm's expenditure is relatively stable. A mineral boom increases exploration activity, which increases the demand for prospectivity maps, which increases the commercial arm's revenue. A mineral bust decreases exploration activity, which decreases demand, which decreases revenue. The public-service arm's expenditure — the flood modelling, the crop assessments, the infrastructure monitoring — does not fluctuate with the commodity cycle, because the flood hazard, the crop conditions, and the infrastructure risks do not diminish when mineral prices fall. Without a smoothing mechanism, the public-service arm's funding would be subject to the commercial arm's volatility, and a mineral downturn would force the office to reduce its public-service output, with the loss of capacity and institutional knowledge that follows from staff reductions.

The smoothing mechanism is a revolving fund that accumulates the commercial surplus in good years and draws it down in poor years, so that the public-service arm's funding is maintained at a stable level across the commercial cycle. The revolving fund is authorized by statute — the United States Geological Survey's revolving fund is authorized by 43 U.S.C. § 364b — and it is managed by the office manager, who sets the target fund balance (the amount the fund should hold to cover a projected downturn) and the draw-down rules (the conditions under which the fund may be drawn below the target). The target fund balance is calibrated to the commercial cycle's historical volatility: a jurisdiction with a volatile mineral sector requires a larger target balance than a jurisdiction with a stable sector. The draw-down rules specify that the fund may be drawn below the target only when the commercial revenue falls below a specified threshold, and that the draw-down must be replenished when the commercial revenue recovers.

The United States Geological Survey's revolving fund allows the Survey to accumulate revenue in good years and draw it down in poor years. The fund's balance is reported in the Survey's annual financial statements, and the Government Accountability Office examines the fund's management during its annual audit. The Ordnance Survey's trading fund framework includes a similar mechanism: the trading fund can accumulate reserves in good years and draw them down in poor years, subject to the Treasury's approval. Without the smoothing mechanism, the revenue-neutral model is fragile. A mineral downturn would force the office to reduce its public-service output, with the loss of capacity that follows. The smoothing mechanism is the financial architecture that makes the revenue-neutral model durable.

5.5 The Audit Regime

The national audit authority examines the commercial arm's accounts, the public-service arm's accounts, and the office's compliance with the firewall's four mechanisms. The audit is conducted annually, and the audit report is published, so that the firewall's maintenance is a matter of public record. The audit is not a formality; its findings can trigger legislative or administrative action where the firewall has been eroded.

The audit authorities that examine geospatial agencies' commercial activities include the Government Accountability Office (GAO) in the United States, the National Audit Office (NAO) in

the United Kingdom, the Cour des comptes in France, and the Australian National Audit Office (ANAO). The GAO, established in 1921 as the General Accounting Office and renamed the Government Accountability Office in 2004, is an independent, nonpartisan agency that works for the United States Congress. The GAO's audits of the United States Geological Survey have examined the Survey's revolving fund management, its commercial activities, and its compliance with the statutory framework. The NAO, established in 1983, audits the Ordnance Survey's trading fund and reports to Parliament on the trading fund's financial performance and its compliance with the Trading Funds Act 1973. The Cour des comptes, established in 1807 by Napoleon, audits IGN France's commercial activities and reports to the President and Parliament. The ANAO, established in 1997, audits Geoscience Australia's commercial activities and reports to Parliament.

The audit examines four questions. Are the commercial arm's revenue and the public-service arm's expenditure accounted for separately, with no commingling of funds? The auditor examines the charge code records, the account structures, and the financial reports to verify that the commercial and public-service funds are separated at every stage of the financial process. Do commercial products carry classification markings that restrict them to the customer and the office, and are the markings enforced in the archive and the distribution system? The auditor examines a sample of commercial products in the archive, verifies that the classification markings are applied, and checks the access log to verify that access is restricted to authorized personnel. Is the office manager the sole authority for commercial releases, and has the authority been delegated or escalated in any case? The auditor examines the commercial release log, verifies that every release is authorized by the office manager personally, and checks whether any release was authorized by another officer (which would constitute a breach of the authority mechanism). Is the revolving fund maintained, and has the fund been used to sustain the public-service arm through the commercial cycle? The auditor examines the fund's balance, its target, its draw-down history, and its replenishment record.

The GAO's audits of the United States Geological Survey's commercial activities have identified issues with the financial architecture and the firewall's maintenance, and the findings have triggered corrective action. A 2015 GAO report (GAO-15-680) examined the Survey's enterprise activities and found that the Survey needed to improve its cost accounting for commercial activities and its tracking of revolving fund balances. The Survey implemented the GAO's recommendations, and subsequent audits have verified the improvements. The NAO's audits of the Ordnance Survey's trading fund have identified similar issues, with similar consequences. The NAO's reports are published on its website and are available to the public, and the reports are reviewed by Parliament's Public Accounts Committee, which may hold hearings on the findings.

The audit's independence from the office is the condition of its credibility. The audit authority reports to the legislature, not to the executive, and the audit authority's budget is appropriated by the legislature, not by the executive. This institutional independence ensures that the audit authority cannot be pressured by the office's parent ministry to soften its findings or to omit issues that might embarrass the ministry. An office that maintains the firewall welcomes the audit, because the audit's clean findings are the evidence the office needs to demonstrate its compliance to its consumers and to the public. An office that has eroded the firewall dreads the audit, because the audit's adverse findings are the trigger for legislative or administrative action that may restructure the office, replace its leadership, or revoke its commercial authority.

5.6 The Stress Tests

The firewall is stress-tested by three predictable pressures. Each pressure tests a different mechanism of the firewall, and each pressure arises from a different feature of the revenue-neutral model.

The first stress test is the commercial downturn. The commercial arm's revenue falls — a mineral price decline reduces exploration activity, which reduces the demand for prospectivity maps, which reduces the commercial revenue. The public-service arm's funding is threatened, because the commercial surplus that supplements the appropriated funding is no longer sufficient. The pressure is to reduce the public-service output — to produce fewer flood risk assessments, to reduce the crop condition bulletin's frequency, to defer the infrastructure monitoring — or to cross-subsidize the public-service arm from the commercial arm's reserves in ways that violate the financial separation. The firewall's response is the revolving fund, which sustains the public-service arm through the downturn by drawing down the accumulated surplus from the good years. The statutory basis prevents the cross-subsidy from masking the commercial arm's losses, because the statute requires the commercial arm's revenue and costs to be reported separately, and a cross-subsidy would be visible in the financial reports and would trigger an audit finding.

The second stress test is the customer who is also a government official. A senior official in a consumer ministry — perhaps the deputy minister of the interior, to whom the office reports — requests a commercial mineral prospectivity product for a specific area of interest. The official's ministry is the office's parent ministry, and the official has authority over the office's budget. The official argues that the product should be provided as a government request, free of charge, rather than as a commercial commission. The pressure is to treat the request as a government request rather than a commercial request, bypassing the commercial release procedure and the commercial charge code. The firewall's response is the authority mechanism, which reserves the commercial release to the office manager, and the classification mechanism, which requires the commercial product to be marked and handled as a commercial product regardless of the customer's government position. The office manager must decline the request as a government request, explain that the product is a commercial product and must be commissioned and paid for as such, and document the interaction in the commercial release log. The documentation is the audit trail that the national audit authority will examine.

The third stress test is political pressure on the public-service arm's intelligence. A minister or a senior official pressures the office to produce intelligence that supports a preferred policy outcome — a flood risk assessment that minimizes the hazard in an area slated for development, or a crop condition assessment that supports a decision to release grain from the strategic reserve. The pressure may be applied directly (a request from the minister to the office manager) or indirectly (a request from the desk chief, who has been instructed by the deputy minister). The firewall's response is the institutional separation, which ensures that the commercial arm's intelligence is not available to the political pressure (the commercial arm's products are confidential and restricted to the customer), and the public-service arm's independence, which is maintained through the office's enabling instrument and the discipline boundaries treated in Chapter 3. An office that passes this stress test produces intelligence that is honest regardless of the pressure. An office that fails it produces intelligence that

is distorted by the pressure, and the distortion is the beginning of the erosion that the firewall is designed to prevent.

5.7 The Firewall in Practice

The firewall is maintained daily by every officer in the office. The data pipeline operator records the commercial arm's data acquisition in the commercial charge code, not the public-service charge code. When the resource economics specialist requests commercial imagery for a mineral prospectivity commission, the operator acquires the imagery under the commercial charge code, records the license terms in the commercial license record, and ingests the data into the processing engine with the commercial charge code attached. The processing engine processes the commercial imagery through the same processing chain as the public-service imagery, but the output is tagged with the commercial charge code, so that the processing output is not commingled with the public-service processing outputs.

The production operator marks commercial products with the commercial classification marking (Confidential, restricted to the customer and the office), not the public-service marking. The production operator archives commercial products in the commercial section of the catalog, with access restricted to the office manager and the resource economics specialist. The production operator does not distribute commercial products through the public-service distribution channel; the commercial distribution channel is separate, and the separation is enforced by the catalog and archive submodule's access controls.

The specialists produce commercial intelligence and public-service intelligence in separate model runs, with separate parameters. The resource economics specialist configures the prospectivity model with commercial data sources (commercial geological data, commercial geophysical data) for commercial commissions, and with public-service data sources (public geological surveys, public geophysical surveys) for public-service standing products. The two sets of model runs are stored in separate project directories, and the commercial model parameters are not available to the public-service specialists. The specialist who produces both commercial and public-service intelligence — the resource economics specialist, who is the only officer besides the office manager who handles commercial products — maintains the separation through the charge code, the data sources, and the model parameters.

The technical lead configures the modeling workbench so that commercial model runs and public-service model runs are separate. The workbench's project management system maintains separate project directories for commercial and public-service work, and the directory permissions restrict access to the commercial directories to the resource economics specialist and the office manager. The technical lead validates commercial model runs separately from public-service model runs, and the validation results are recorded under the commercial charge code.

The office manager authorizes every commercial release personally. The manager reviews the commercial product for firewall compliance (verifying that the product was produced from commercial data sources and commercial model runs, not from public-service data or public-service model runs), verifies the classification marking (Confidential, restricted to the customer and the office), verifies the charge code (commercial, not public-service), verifies the customer identity and the agreed price against the commercial customer register, and signs the commercial release log. The

manager refuses to delegate the authority, even under pressure from the desk chief or the deputy minister. The refusal is not a matter of personal preference; it is a matter of statutory obligation, and the manager who delegates the authority has breached the firewall and is subject to the consequences that the statute prescribes. The manager reports any pressure to erode the firewall to the national audit authority, as a matter of the audit mechanism's requirements.

5.8 Worked Example: A Commercial Commission Under Stress

A mining company — Meridian Resources — commissions a mineral prospectivity assessment for a 500-square-kilometre area in the northern highlands. The commission is for a custom prospectivity map integrating geological, aeromagnetic, and geochemical data, with a delivery timeline of eight weeks and an agreed price of $45,000. The office manager accepts the commission, assigns it to the resource economics specialist, and opens a commercial charge code (COMM-2024-017) for the work.

The specialist configures the prospectivity model with commercial data sources: commercial aeromagnetic data purchased from a geophysical survey company, commercial geochemical data purchased from a laboratory, and the public geological map (which is freely available). The specialist runs the model under the commercial charge code, validates the results against known deposits in the area, and produces the prospectivity map. The production operator applies the Confidential classification marking (restricted to Meridian Resources and the office), converts the map to the customer's required format (GeoTIFF and PDF report), and archives the product under the commercial charge code.

Four weeks into the commission, the deputy minister of the interior contacts the office manager. The deputy minister explains that the government is considering a land-use policy change in the northern highlands that would affect mineral exploration, and that the government would benefit from seeing the prospectivity assessment that the office is producing for Meridian Resources. The deputy minister requests that the prospectivity assessment be shared with the government as an internal government product, rather than as a commercial product restricted to Meridian Resources.

The office manager faces the second stress test: the customer who is also a government official, complicated by the fact that the official is the deputy minister to whom the office reports. The manager must decide how to respond.

The manager's response is governed by the firewall's authority mechanism and classification mechanism. The prospectivity assessment is a commercial product, commissioned by Meridian Resources under a commercial contract, produced under a commercial charge code, and marked Confidential (restricted to Meridian Resources and the office). The product cannot be shared with the government as an internal government product without Meridian Resources's consent, because the product is Meridian Resources's property under the commercial contract. The manager explains this to the deputy minister: the product is a commercial commission, the product is the customer's property, and the office cannot release it to a third party (including the government) without the customer's consent. The manager offers an alternative: the office can produce a separate public-service prospectivity assessment for the northern highlands, using public-service data sources and public-service model parameters, and deliver it as an internal government product. The public-service

assessment would not use Meridian Resources's commercial data, would not be the same product as the commercial commission, and would be produced under a public-service charge code.

The deputy minister accepts the alternative. The office manager assigns the resource economics specialist to produce a public-service prospectivity assessment for the northern highlands, using public geological map data, public aeromagnetic data (from the national geological survey's open-file data), and public geochemical data (from the national geochemical atlas). The specialist configures a new model run under a public-service charge code (PUB-2024-042), produces the assessment, and delivers it to the deputy minister as an internal government product.

The worked example illustrates the firewall's operation under stress. The commercial product (COMM-2024-017) was produced under the commercial charge code, with commercial data sources, and marked Confidential. The public-service product (PUB-2024-042) was produced under the public-service charge code, with public-service data sources, and marked Internal. The two products were produced from different data sources, under different charge codes, with different classification markings, and through different distribution channels. The firewall's four mechanisms — financial (separate charge codes), classification (Confidential vs Internal), authority (the manager's decision), and audit (both products are documented in the catalog and archive, and the manager's response to the deputy minister's request is documented in the commercial release log) — operated as designed. The commercial customer's product was protected. The government's intelligence needs were met through a separate public-service product. The firewall was maintained.

Chapter 6 — The Acquisition-to-Archive Pipeline

Part II — The Geospatial Intelligence Method: The acquisition-to-archive pipeline as the shared backbone of the office's work, comprising six capabilities with a single shared stack and swappable modeling heads.

6.1 The Pipeline Concept

The acquisition-to-archive pipeline is the shared backbone of the geospatial office's work. It comprises six capabilities: acquisition, base layer, processing, modeling, production, and archive. The first five capabilities are identical across all domain applications; only the modeling capability is domain-specific. Adopting the pipeline model is the single most consequential design decision in the office's establishment. An office that adopts the pipeline produces consistent, compatible, and auditable products across all its domains. An office that does not produces inconsistent, incompatible, and unauditable products, regardless of the quality of its individual analysts.

The pipeline concept is drawn from the operational practice of national mapping agencies and Earth observation programmes that have adopted shared-backbone architectures. The United States Geological Survey's National Map programme, established in 2001 as a consolidation of the topographic, hydrographic, and land-cover mapping functions, maintains a single processing pipeline that produces the topographic base, the elevation layer (the 3D Elevation Program, or 3DEP), the hydrography layer (the National Hydrography Dataset, or NHD), and the land cover layer (the National Land Cover Database, or NLCD). Domain-specific analysis — hazard assessment, mineral resource evaluation, environmental monitoring — is conducted on top of the shared base, using the same data through different analytical methods. The consolidation eliminated the redundant data acquisition and processing that had previously been conducted separately by each programme.

The European Space Agency's Copernicus programme, operational since 2014, maintains a single Sentinel satellite processing pipeline — the Copernicus Ground Segment — that serves all six Copernicus services: the Emergency Management Service, the Land Monitoring Service, the Marine Environment Monitoring Service, the Atmosphere Monitoring Service, the Security Service, and the Climate Change Service. Each service consumes the processed Sentinel data through a common interface and applies its own analytical methods, but the data acquisition, processing, and distribution are conducted through the shared ground segment. The shared ground segment was a deliberate design choice: the European Commission estimated that the shared infrastructure reduced the programme's overall cost by approximately 30 per cent compared to six separate satellite-data pipelines.

Geoscience Australia's Earth observation data pipeline, maintained through the agency's Spatial and Earth Monitoring Division, serves the agency's mineral, environmental, and hazard programs from a single acquisition and processing backbone. The agency's satellite data reception facility in Alice Springs acquires data from Landsat, Sentinel, and other satellites, and the acquired data is processed through a single processing chain before being distributed to the agency's analytical programmes. The shared backbone ensures that the agency's mineral prospectivity assessments, its earthquake hazard maps, and its environmental monitoring products are all based on the same processed data, with the same georeferencing, the same calibration, and the same base layers.

The pipeline model is the common architecture of the national geospatial agencies that produce consistent, high-quality output at scale. Agencies that have not adopted the model — where each programme acquires its own data, maintains its own base layers, and runs its own processing chain — produce inconsistent output, because the data, the processing, and the base layers differ across programmes even within the same agency. The inconsistency is not always visible to the consumer, because the products look similar, but the inconsistency undermines the products' credibility when the consumer attempts to cross-reference them.

6.2 The Six Capabilities

Every geospatial product passes through six capabilities, from the moment data is acquired to the moment the product is archived. The capabilities are sequential in the sense that a product passes through them in order, but they are concurrent in operation, because the office operates all six simultaneously, with different products at different stages of the pipeline at any given time.

Acquisition covers the tasking of satellites, the purchasing of imagery, the ingestion of sensor feeds, and the management of data licenses. The acquisition capability is the entry point of the pipeline, and its quality determines the quality of every subsequent product. The acquisition capability is treated in detail in Chapter 7.

The base layer covers the maintained set of base layers — the digital elevation model, land cover, hydrography, administrative boundaries, and the cadastre feed — that provide the geographic context for every analysis. The base layer capability is the geographic foundation of the pipeline, and its integrity is the non-negotiable precondition for every product. The base layer capability is treated in detail in Chapter 8.

Processing covers the transformation of raw acquired data into analysis-ready products through georeferencing, calibration, mosaicking, projection management, cloud masking, and data fusion. The processing capability is the technical core of the non-analytical pipeline, and its quality determines whether the analysts can work with the data at all. The processing capability is treated in detail in Chapter 9.

Modeling covers the domain-specific analytical layer, conducted through the modeling workbench with its domain plug-ins. The modeling capability is the only capability that varies by domain, and it is the capability that produces the office's intelligence. The modeling capability is treated in detail in Chapter 10.

Production covers the transformation of model outputs into final deliverables through cartographic design, symbology application, classification marking, and format conversion. The production capability is the communication stage of the pipeline, and its quality determines whether the intelligence reaches the consumer in a usable form. The production capability is treated in detail in Chapter 11.

Archive covers the preservation of the office's outputs and the maintenance of the catalog that makes them retrievable. The archive capability is the institutional memory of the pipeline, and its quality determines whether the office can reconstruct, audit, and learn from its past work. The archive capability is treated in detail in Chapter 12.

The first five capabilities are identical across all domain applications. The acquisition capability acquires data for flood models, crop assessments, infrastructure monitoring, and mineral

prospectivity through the same satellite tasking, imagery purchasing, and sensor feed ingestion processes. The base layer capability maintains the same elevation, land cover, hydrography, and cadastre layers for all four domains. The processing capability georeferences, calibrates, and mosaics the imagery through the same processing chain regardless of the intended domain application. The production capability applies the same cartographic standards, classification markings, and format conversions to all products. The archive capability maintains the same provenance chain, version history, and access logging for all products. Only the modeling capability varies by domain, because only the analytical method is domain-specific.

6.3 The Shared-Backbone Principle

The shared-backbone principle holds that the five non-domain capabilities are conducted through a single, shared stack, not through four parallel stacks. This principle prevents the domain-silo failure, which is the most common failure mode in government GIS deployments.

The domain-silo failure occurs when each domain acquires its own data, maintains its own base layers, runs its own processing chain, and produces its own products, with no shared infrastructure. The flood risk team purchases its own satellite imagery, maintains its own elevation model, runs its own georeferencing chain, and archives its products in its own folder. The food security team does the same, with different imagery, a different elevation model (or the same model at a different version), a different georeferencing chain, and a different archive folder. The result is four incompatible data silos, four redundant processing chains, four sets of base layers that diverge over time (because each team updates its base layers independently, and the updates are not synchronized), and four archives that cannot be cross-referenced (because the products are in different formats, with different projections, and different metadata standards).

The domain-silo failure defeats the purpose of having a single geospatial office, because the office's products are no more compatible than the products of four separate offices. A flood risk assessment produced by the flood team and a crop condition assessment produced by the food security team may use different base layers, different projections, and different processing parameters, and the two assessments cannot be overlaid or cross-referenced without extensive and error-prone manual reconciliation. The incompatibility is not visible to the consumer until the consumer attempts to use the products together, at which point the incompatibility becomes apparent and the consumer's trust in the office's products is undermined.

The domain-silo failure has been documented in multiple government geospatial agencies. The United States federal government's geospatial activities were fragmented across agencies for decades before the National Spatial Data Infrastructure policy (Executive Order 12906, 1994) began the consolidation. A 2004 National Research Council report, "Geographic Information for the 21st Century," found that the federal government's geospatial data activities were characterized by "duplication, redundancy, and lack of coordination" and estimated the annual cost of the redundancy at hundreds of millions of dollars. The European Union's INSPIRE Directive (2007/2/EC) was motivated by similar findings: a 2004 European Commission study estimated that the fragmentation of geospatial data across member states cost the European economy approximately €1 billion annually in lost productivity and redundant data acquisition.

The shared-backbone principle is maintained through the office's information system architecture, treated in Chapter 29. The architecture comprises a single acquisition desk that serves all domains, a single spatial data infrastructure that maintains the base layers for all domains, a single processing engine that processes imagery for all domains, a single modeling workbench with domain plug-ins that hosts the analytical models, a single product studio that produces the final deliverables for all domains, and a single catalog and archive that preserves all products. An office that adopts the principle without the architecture, or the architecture without the principle, will not achieve the pipeline's benefits. The principle and the architecture are inseparable: the principle is the doctrine, the architecture is the implementation, and both must be adopted together.

The United States Geological Survey's enterprise GIS architecture maintains a single data pipeline that serves all the Survey's programs. The National Map, the 3D Elevation Program, the National Hydrography Dataset, and the National Land Cover Database are all maintained through the same acquisition, processing, and distribution infrastructure. Program-specific analysis — the earthquake hazard assessments produced by the Earthquake Hazards Program, the mineral resource assessments produced by the Mineral Resources Program, the water resource assessments produced by the Water Resources Program — is conducted on top of the shared base. The Copernicus programme's ground segment architecture maintains a single processing chain for all Sentinel data, with the Copernicus services consuming the processed data through a common interface. Geoscience Australia's spatial information infrastructure maintains a single data pipeline for all the agency's programs. Agencies that maintain the shared backbone produce consistent, compatible products; agencies that do not produce inconsistent, incompatible products, and the difference is visible in both the quality and the cost of the output.

6.4 The Swappable Modeling Heads

The modeling capability is the only pipeline capability that varies by domain, and it varies through the mechanism of swappable modeling heads. The modeling workbench is the shared environment that hosts the analytical models. The domain plug-ins are the swappable heads that implement the domain-specific methods.

The hydrology plug-in implements the flood risk models: rainfall-runoff models (such as the Sacramento Soil Moisture Accounting model or the Green-Ampt infiltration model), hydraulic models (such as HEC-RAS 2D, MIKE 21, or TUFLOW), and inundation mapping routines. The plug-in consumes the digital elevation model, the land cover, the hydrography, the rainfall data, and the river stage data, and it produces the flood inundation maps, the exposure reports, and the evacuation guidance that the flood risk domain delivers.

The crop model plug-in implements the food security models: vegetation index calculators (NDVI, NDWI, EVI), crop growth models (such as the World Food Studies (WOFOST) model or the Decision Support System for Agrotechnology Transfer (DSSAT)), and yield forecasting routines. The plug-in consumes the optical imagery, the weather data, the soil moisture data, the cadastre parcels, and the farmer-reported planting data, and it produces the crop condition assessments, the yield forecasts, and the procurement advisories that the food security domain delivers.

The change detection plug-in implements the energy infrastructure models: image differencing algorithms, object-based change detection, and vulnerability scoring routines. The plug-in consumes

the high-resolution imagery, the infrastructure asset registry, the elevation data, and the vegetation cover data, and it produces the infrastructure risk maps, the encroachment alerts, and the maintenance advisories that the energy infrastructure domain delivers.

The prospectivity plug-in implements the minerals models: weights-of-evidence models, fuzzy logic models, and machine learning classifiers (random forest, support vector machine). The plug-in consumes the geological survey data, the aeromagnetic and gravity geophysics, the geochemistry, the hyperspectral imagery, and the historical exploration data, and it produces the prospectivity maps, the resource estimates, and the critical minerals inventory that the minerals domain delivers.

The plug-ins share the workbench's data interface (they all read from the same spatial data infrastructure), its parameter management (they all use the same parameter configuration system), its validation framework (they all report validation results in the same format), and its provenance recording (they all record model version, parameters, run date, and validation results in the same provenance structure). They differ in the analytical methods they implement. The shared interface and the differing methods are the design that allows the pipeline to serve four domains through one workbench rather than four.

The swappable-heads design allows the pipeline to serve multiple domains without becoming a domain-specific pipeline. The alternative designs are the single-domain pipeline, which serves one domain well but cannot serve others, and the parallel-pipeline design, which serves multiple domains through parallel pipelines at the cost of the shared-backbone benefits. The swappable-heads design serves multiple domains through a single pipeline with domain-specific modeling, combining the benefits of the shared backbone with the flexibility of domain-specific analysis. The design is drawn from the practice of national geospatial agencies that maintain shared analytical platforms with domain-specific applications, including the United States Geological Survey's hazard modeling framework (the National Water Model, the National Seismic Hazard Model, the National Volcanic Early Warning System) and the European Space Agency's Copernicus services platform (which hosts the six Copernicus services on a common data and processing infrastructure).

The swappable-heads design has implications for the office's complement. The specialists who operate the modeling plug-ins must be domain experts in their respective tracks, but they share the workbench's configuration, its data interface, and its validation framework with the other specialists and the technical lead. The shared workbench allows the specialists to collaborate, to cross-validate, and to cover for each other during absence or surge. A specialist who works in a domain-specific silo, with no shared workbench, cannot collaborate, cannot cross-validate, and cannot be covered. An office that allows its specialists to work in silos has lost the pipeline's benefits for the modeling capability, even if it maintains the shared backbone for the other five capabilities.

6.5 The Provenance Chain

The provenance chain is the traceable record from acquisition through archive. Every product is traceable to the data it was produced from, the processing parameters that were applied, the model version that was used, the validation that was conducted, the production parameters that were applied, and the archive record that preserves the product. The provenance chain makes the office's products auditable, reproducible, and defensible. An office that maintains the provenance chain can answer the

question "how was this product produced?" for any product it has ever produced. An office that does not cannot, and its products are not auditable.

The provenance chain is established at the acquisition stage and maintained through every subsequent stage. At acquisition, the data source (e.g., Sentinel-2, Landsat 8, commercial Maxar WorldView-3), the license terms (free and open, commercial, restricted), the acquisition date (the date the imagery was acquired by the satellite), the ingestion date (the date the office received the data), and the intended domain consumer (flood risk, food security, energy infrastructure, minerals) are recorded. At processing, the processing parameters (georeferencing method, calibration coefficients, mosaicking parameters, cloud masking threshold), the processing date, the processing software version (e.g., ESA SNAP 9.0, GRASS GIS 8.3), and the processed data's checksum are recorded. At modeling, the model version (e.g., HEC-RAS 2D version 6.3), the model parameters (Manning's n values, boundary conditions, grid resolution), the validation results (the accuracy metrics and the validation method), and the model run date are recorded. At production, the product version, the production date, the classification marking (Public, Internal, Confidential, Security Confidential), the cartographic template version, and the intended recipient are recorded. At archive, the product identifier (a unique identifier that links the product to its provenance record), the retention designation (the period for which the product is retained), and the access permissions (who may access the product) are recorded.

The chain is unbroken. Any break in the chain — a missing data source, an unrecorded processing parameter, an undocumented model version — is a defect that must be corrected before the product is released. The defect is corrected by tracing the product back through the chain to the point where the record was lost and re-creating the record from the available evidence (the processing log, the model run log, the production log). If the record cannot be re-created, the product is rejected, because a product without a complete provenance chain is a product whose history cannot be verified and whose reliability cannot be assessed.

The United States Geological Survey's data management policy, documented in Survey Manual Chapter 502.1, requires that every published product be traceable to its source data, its processing history, and its validation record. The policy is enforced through the Survey's publication review process, which requires the submitting scientist to provide the complete provenance record before the product is approved for publication. The Copernicus programme's data quality framework, documented in the Copernicus Quality Assurance Framework, requires that every Sentinel product be traceable to its acquisition (the satellite, the sensor, the acquisition time), its processing (the processing baseline, the processor version, the processing parameters), and its quality control (the quality flags, the accuracy assessment). Geoscience Australia's data management framework, documented in the agency's Information Management Plan, requires that every product be traceable through its production chain. The provenance chain is the integrity guarantee that the national geospatial agencies provide, and it is the guarantee that the pipeline is designed to establish for the geospatial office.

6.6 The Cadence Design

The pipeline operates on two cadences: a fixed pipeline cadence and a variable modeling cadence. The fixed pipeline cadence is the rhythm at which the acquisition, base layer, processing,

production, and archive capabilities operate. It is fixed because these capabilities are shared across all domains and must operate on a predictable schedule. The variable modeling cadence is the rhythm at which the modeling capability operates. It is variable because the four domains operate on different schedules.

The fixed cadence comprises daily ingestion (the data pipeline operator checks the overnight acquisitions every morning), daily processing queue management (the operator configures and monitors the day's processing runs), weekly base layer integrity checks (the technical lead verifies synchronization with the source authorities every Monday), monthly archive audits (the production operator verifies the archive's completeness and integrity on the first working day of each month), and production on demand (the production operator produces products as the specialists' model outputs become available). The fixed cadence is the office's heartbeat, maintained by the data pipeline operator, the production operator, and the technical lead, with the office manager overseeing.

The variable cadence is determined by each domain's operational requirements. The flood risk domain operates on a continuous cadence: the sensor feeds are ingested continuously, and the flood models are run whenever the sensor data indicates a potential event. During a flood event, the cadence intensifies to event-driven production, with model runs every few hours and products delivered to the emergency management authority within hours. The food security domain operates on a weekly cadence: the crop condition bulletin is produced every week and delivered by close of business each Wednesday. During the growing season, the cadence intensifies at the seasonal milestones (planting, heading, harvest), when the yield forecasts are updated. The energy infrastructure domain operates on a monthly cadence: the infrastructure monitoring cycle is conducted once a month, with the monitoring products delivered by the fifth business day. The cadence intensifies when an incident is detected or reported, triggering an event-driven response. The minerals domain operates on a project-based cadence: commercial commissions are produced on the timeline agreed per project (typically 4-12 weeks), and the annual standing products (the national geological map series, the critical minerals inventory) are produced on an annual cycle.

Decoupling the fixed pipeline cadence from the variable modeling cadence is the design move that allows the pipeline to serve four domains with incompatible schedules, and it is the move that makes a 6-person office viable. Without the decoupling, the office would need to synchronize the four domains' schedules, which is impossible because the flood event does not wait for the weekly crop bulletin's deadline, and the commercial commission's delivery timeline does not align with the monthly monitoring cycle. The decoupling allows each domain to pull from the pipeline on its own schedule, while the pipeline operates on a fixed schedule that serves all domains.

The cadence design is drawn from the operational practice of newsrooms and intelligence agencies, which maintain a fixed production cadence for the shared infrastructure and a variable publication cadence for the domain-specific output. A newsroom maintains a fixed copy desk, editing desk, and production desk, while reporters and section editors operate on variable cadences determined by their beats. The copy desk does not wait for the sports reporter to file before it processes the political reporter's copy; each reporter files on their own schedule, and the copy desk processes whatever comes in. An intelligence agency maintains a fixed collection and processing cadence and a variable analysis cadence determined by intelligence priorities. The geospatial office follows the same pattern: the pipeline capabilities operate on a fixed cadence (the data pipeline

operator processes whatever data arrives, regardless of which domain will use it), and the modeling capability operates on a variable cadence (the specialists pull from the processed data on their own schedules). The decoupling allows the small office to serve multiple domains without collapsing under schedule conflicts.

6.7 The Pipeline and the Domains

The pipeline and the domains are related through the modeling capability. The pipeline provides the shared backbone; the domains provide the swappable modeling heads. The pipeline is treated in Part II; the domains are treated in Part IV; the roles that operate the pipeline and the models are treated in Part III. The three parts are interdependent, and the pipeline is the architecture that connects them.

The pipeline is the subject of the office's daily work, treated in Part III. The data pipeline operator operates the acquisition, processing, and base layer capabilities. The production operator operates the production and archive capabilities. The technical lead configures the pipeline and validates the models. The specialists operate the modeling capability. The office manager oversees the pipeline and authorizes releases. The daily work of each seat is the operation of a part of the pipeline, and the pipeline is the architecture that makes the daily work coherent. The Job Breakdown Sheets in Part III are the operational specifications of the pipeline capabilities.

The pipeline is the subject of the office's information system, treated in Chapter 29. The six submodules of the geospatial information system correspond to the six capabilities of the pipeline: the spatial data infrastructure corresponds to the base layer capability, the acquisition desk corresponds to the acquisition capability, the processing engine corresponds to the processing capability, the modeling workbench corresponds to the modeling capability, the product studio corresponds to the production capability, and the catalog and archive corresponds to the archive capability. The system is the technical implementation of the pipeline architecture. The pipeline is the doctrine; the information system is the implementation; the daily work is the operation. An office that maintains the alignment of doctrine, implementation, and operation produces consistent, high-quality intelligence. An office that allows them to diverge produces inconsistent, low-quality output.

6.8 Worked Example: A Day in the Pipeline

The pipeline's operation is best understood through a single day's work. On a Tuesday morning in the growing season, the office's pipeline processes products for three domains simultaneously, with the fourth domain (minerals) in a quiet period between commissions.

At 07:30, the data pipeline operator arrives and begins the morning ingestion check (section 15.3). The overnight acquisition log shows that a Sentinel-2 image was acquired over the agricultural region at 06:00 UTC, with 15 per cent cloud cover — suitable for the crop condition assessment. A Landsat 8 image was acquired over the flood basin at 04:30 UTC, with 60 per cent cloud cover — too cloudy for optical analysis but useful for the thermal band, which can identify saturated soil. The sensor feeds from the 15 rain gauges and 8 river stage sensors are all active and delivering current data; the river stages are within normal range, and no flood event is indicated.

The operator ingests the Sentinel-2 image into the processing engine under the public-service charge code (PUB-2024-118), records the license terms (Copernicus free and open data, commercial use permitted with attribution), and configures the processing run: georeferencing using the

Sentinel-2 Level-2A processing baseline (which includes atmospheric correction and surface reflectance calculation), mosaicking not required (single scene), cloud masking using the Scene Classification Layer (SCL) at confidence threshold 60 per cent. The operator submits the processing job and moves to the next task.

At 08:00, the office manager conducts the morning production review (section 13.3) with the data pipeline operator and the production operator. The review covers the overnight log, the day's processing queue, the modeling queue, the production queue, and the flood event status. The manager confirms that the day's priorities are: (1) the weekly crop condition bulletin, due Wednesday by close of business; (2) the monthly energy infrastructure monitoring cycle, due Friday by the fifth business day; (3) the ongoing processing of the Sentinel-2 image for the crop bulletin. No flood event is indicated.

At 08:30, the processing engine completes the Sentinel-2 processing run. The technical lead reviews the processing output (section 14.3): the geometric accuracy is within the Sentinel-2 specification (RMSE < 10.3 metres at the 95 per cent confidence level), the radiometric calibration is consistent with the previous acquisition, the cloud masking has identified 15 per cent cloud cover (matching the acquisition log), and the provenance record is complete. The technical lead releases the processed data to the modeling workbench.

At 09:00, the resource economics specialist begins the crop condition assessment for the weekly bulletin (section 18.3). The specialist configures the crop model plug-in with the parameters for the current growth stage (the crops are in the heading stage), retrieves the processed Sentinel-2 image from the workbench, calculates the NDVI and NDWI, stratifies the indices by parcel using the cadastre feed, and produces the parcel-level condition assessments. The specialist updates the yield forecast by integrating the current condition assessment with the historical yield data and the seasonal progress. The specialist validates the assessment against the ground truth observations from the agricultural extension service (3 of 4 ground truth sites confirm the assessment; the 4th site shows a 10 per cent discrepancy that the specialist attributes to a local irrigation difference). The specialist documents the model run, the validation, and the discrepancy in the provenance record.

At 10:00, the hazard specialist begins the monthly energy infrastructure monitoring cycle (section 17.4). The specialist retrieves the high-resolution imagery acquired for the monitoring cycle (Maxar WorldView-3, 0.31-metre resolution, acquired the previous week under the public-service charge code PUB-2024-115), configures the change detection plug-in with the standard parameters, and runs the change detection against the previous month's imagery. The detection identifies 7 changes: 3 vegetation encroachment alerts, 2 third-party construction alerts, 1 ground subsidence indicator, and 1 false alarm (a shadow change). The specialist reviews each change, cross-references the changes with the infrastructure asset registry, and produces the vulnerability scores. The specialist generates the infrastructure risk map and the encroachment alerts with Security Confidential classification.

At 13:00, the production operator begins the production runs for both domains (section 16.3). For the crop condition bulletin, the operator applies the cartographic template, the symbology, and the split classification marking (Public for the summary, Internal for the detailed assessment). For the infrastructure monitoring products, the operator applies the Security Confidential marking. The operator converts both products to the consumers' required formats (GeoTIFF and PDF for the crop

bulletin; GeoTIFF and Shapefile for the infrastructure monitoring) and submits the draft products to the office manager for review.

At 14:30, the office manager reviews both products (section 13.3). The manager verifies the classification markings (Public/Internal split for the crop bulletin; Security Confidential for the infrastructure monitoring), confirms that the provenance records are complete, and authorizes release. The crop condition bulletin is scheduled for delivery on Wednesday morning. The infrastructure monitoring products are delivered through the secure data interface to the energy sector entity and the sectoral regulator.

At 16:00, the production operator archives both products in the catalog and archive submodule (section 16.5). The operator records the product identifiers, the provenance links (tracing the crop bulletin to the Sentinel-2 image, the processing run, the crop model run, and the production run; tracing the infrastructure monitoring to the WorldView-3 image, its processing run, the change detection run, and the production run), the retention designations (permanent for the crop bulletin; 10 years for the infrastructure monitoring), and the access permissions (Public/Internal for the crop bulletin; Security Confidential for the infrastructure monitoring).

At 16:30, the production operator conducts the daily classified access log review (section 16.6). The review covers the day's access to the Security Confidential products in the archive. All access is by authorized personnel (the hazard specialist, the production operator, and the office manager), and no unusual access patterns are identified.

At 17:00, the office manager conducts the day-end production review. The day's priorities have been addressed: the crop condition bulletin is on schedule for Wednesday delivery, the monthly infrastructure monitoring cycle is complete, and no flood event has been indicated. The pipeline has processed data for two domains, produced products for two domains, and archived both, all in a single day, through a shared backbone that served both domains without duplication or conflict. This single day illustrates the pipeline's core value: one infrastructure, four domains, one office, consistent output.

Chapter 7 — Acquisition and Ingest

7.1 The Acquisition Function

Acquisition and ingest is the first capability of the pipeline. It covers the tasking of satellites, the purchasing of imagery, the ingestion of sensor feeds, and the management of data licenses. The capability is the entry point of the pipeline, and its quality determines the quality of every subsequent product. An office that acquires the right data, at the right resolution, at the right time, with the right license, produces intelligence that is timely, accurate, and legally usable. An office that acquires the wrong data, or the right data at the wrong time, or the right data with the wrong license, produces intelligence that is delayed, inaccurate, or legally compromised.

The acquisition function is conducted through the acquisition desk submodule of the geospatial information system, treated in Chapter 29. The acquisition desk manages satellite tasking requests, imagery purchase orders, sensor feed subscriptions, and data license records. It provides the interface through which specialists request data, the data pipeline operator acquires it, the technical lead configures the acquisition parameters, and the office manager authorizes purchases above the threshold. The acquisition desk's design and use are specified in the job breakdown sheets in Chapter 15.

The acquisition function has four sub-capabilities, each treated in a separate section of this chapter. Satellite tasking (section 7.2) covers the process of requesting that a satellite acquire imagery of a specific area at a specific time. Imagery purchasing (section 7.3) covers the process of buying imagery from commercial providers, either as tasking requests or as archive imagery. Sensor feed ingestion (section 7.4) covers the process of receiving and processing continuous data streams from ground-based sensors. License management (section 7.5) covers the process of recording, tracking, and enforcing the data license terms that govern the use of the data the office acquires.

7.2 Satellite Tasking

Satellite tasking is the process of requesting that a satellite acquire imagery of a specific area at a specific time. The capability is available through commercial imagery providers and through national Earth observation programmes where the office's government operates satellites or has access to allied satellites.

The commercial high-resolution satellite market is dominated by three providers. Maxar Technologies (United States), through its WorldView constellation (WorldView-1, launched 2007, 0.5m panchromatic; WorldView-2, launched 2009, 0.46m panchromatic and 8-band multispectral; WorldView-3, launched 2014, 0.31m panchromatic and 8-band multispectral plus 8-band shortwave infrared; WorldView-4, launched 2016, 0.31m panchromatic and 4-band multispectral), provides the highest-resolution commercial satellite imagery available. Airbus Defence and Space (Europe), through its Pleiades constellation (Pleiades-1A, launched 2011, 0.5m panchromatic and 2m multispectral; Pleiades-1B, launched 2012; Pleiades Neo, launched 2021-2022, 0.3m panchromatic and 1.2m multispectral), provides high-resolution imagery with a daily revisit capability. Planet Labs (United States), through its Dove constellation (approximately 200 CubeSats providing 3m multispectral imagery with daily global coverage) and its SkySat constellation (21 satellites

providing 0.5m panchromatic and 2m multispectral imagery), provides high-frequency, moderate-resolution imagery that complements the higher-resolution but less frequent Maxar and Airbus imagery.

The tasking request specifies five parameters. The area of interest is the geographic area to be imaged, specified as a bounding box, a polygon, or a point with a radius. The required resolution is the minimum ground sample distance (GSD) that the application requires: flood inundation mapping typically requires 10m or better (Sentinel-2); crop condition assessment requires 10-30m (Sentinel-2 or Landsat); infrastructure change detection requires 0.5m or better (Maxar or Airbus); mineral prospectivity may require 0.5m for structural mapping but 30m for regional geological context. The required spectral bands are the bands the application needs: panchromatic for high-resolution mapping, multispectral for vegetation analysis, shortwave infrared for mineral mapping, thermal infrared for soil moisture and infrastructure monitoring. The required acquisition window is the time period within which the imagery must be acquired: a flood event may require imagery within 24 hours; a crop condition assessment may require imagery within a week; an infrastructure monitoring cycle may require imagery within a month. The maximum acceptable cloud cover is the percentage of the image that may be obscured by clouds: 0 per cent for infrastructure monitoring (clouds hide the infrastructure); 20 per cent for crop condition assessment (some cloud cover is acceptable because the assessment is based on clear-sky pixels); 100 per cent for flood extent mapping using radar (radar penetrates clouds).

The provider confirms or declines the request based on the satellite's orbit (the revisit time for the area of interest), competing demand (other customers may have tasked the same satellite for the same area), and weather conditions (for optical satellites, cloud cover may prevent acquisition). The provider's confirmation includes the expected acquisition date, the expected image quality, and the price. The provider's decline includes the reason (orbit conflict, weather, or capacity) and the next available acquisition opportunity.

The United States government operates the Landsat programme (Landsat 8, launched 2013, 30m multispectral and 15m panchromatic, 16-day revisit; Landsat 9, launched 2021, same specifications), which provides moderate-resolution imagery on a regular revisit cycle without tasking. The government tasks commercial satellites through the National Geospatial-Intelligence Agency's EnhancedView programme, a $7.3 billion contract awarded in 2010 to Maxar (then DigitalGlobe) and GeoEye (later merged into Maxar), which provides the NGA and authorized government users with access to commercial high-resolution imagery. The European Space Agency operates the Sentinel programme (Sentinel-2A, launched 2015, 10m/20m/60m multispectral, 5-day revisit with Sentinel-2B; Sentinel-2B, launched 2017), which provides moderate-resolution imagery on a regular revisit cycle without tasking, and tasks commercial satellites through the Copernicus programme's Contributing Missions for high-resolution imagery.

The geospatial office follows the same model. The office subscribes to free and open moderate-resolution imagery (Sentinel-2 and Landsat 8/9) for routine monitoring in the flood risk, food security, and energy infrastructure domains. The office tasks commercial high-resolution imagery (Maxar WorldView-3, Airbus Pleiades Neo) for specific applications — post-event flood assessment (where the moderate-resolution imagery is insufficient to identify building-level damage), infrastructure change detection (where the moderate-resolution imagery is insufficient to identify

encroachment on transmission corridors), and mineral prospectivity (where the moderate-resolution imagery is insufficient for structural geological mapping) — where moderate-resolution imagery is insufficient. Tasking is conducted through the acquisition desk: specialists submit tasking requests, the data pipeline operator submits them to the provider, the technical lead reviews the tasking parameters, and the office manager authorizes purchases above the threshold.

7.3 Imagery Purchasing

Imagery purchasing is the process of buying imagery from commercial providers, either as tasking requests (new acquisitions) or as archive imagery (previously acquired imagery from the provider's archive). The capability is governed by three constraints: the office's data acquisition budget, the procurement framework, and the data license terms.

The budget determines how much the office can spend. The office's data acquisition budget is part of the office's overall budget, and it is managed by the office manager. The budget must cover both the routine acquisitions (sensor feed subscriptions, standing satellite subscriptions) and the ad hoc acquisitions (commercial tasking for specific applications). The budget is managed through the acquisition desk, which tracks the cumulative spending against the budget and alerts the office manager when the spending approaches the budget limit. The budget management is the financial discipline that prevents the office from overspending on data acquisition and underinvesting in the analytical work that the data supports.

The procurement framework determines how the office can spend the budget. The framework is established by the government's procurement authority and provides the contracting vehicles through which the office purchases commercial imagery. The most common vehicle is the framework agreement: the procurement authority negotiates a master agreement with each commercial provider that establishes the pricing, the license terms, the delivery timelines, and the quality standards, and the office purchases under the master agreement without negotiating individual contracts. The framework agreement reduces the transaction costs of imagery purchasing (the office does not need to negotiate terms for each purchase) and ensures that all government agencies operate under consistent terms (which simplifies inter-agency data sharing).

The United States government purchases commercial imagery through the National Geospatial-Intelligence Agency's contracting vehicles. The EnhancedView programme, awarded in 2010, is the principal vehicle for commercial high-resolution satellite imagery. The GSA Schedule (the General Services Administration's Federal Supply Schedule) provides another vehicle for commercial geospatial products and services. The European Space Agency purchases commercial imagery through the Copernicus programme's contracting framework, which provides framework agreements with commercial providers for the Copernicus Contributing Missions. The model is consistent across comparative practice: the government negotiates framework agreements that provide standardized terms, and the agencies purchase under the frameworks.

The license terms determine what the office can do with the imagery it buys. The license terms are the legal contract between the provider and the government, and they specify the permitted uses, the prohibited uses, the redistribution rights, the attribution requirements, and the duration of the license. The license terms for commercial imagery are typically more restrictive than the license terms for free and open data: commercial imagery may be licensed for internal government use only,

prohibiting redistribution to the public; commercial imagery may be licensed for a specific project, prohibiting use in other projects; commercial imagery may be licensed for a specific number of users, prohibiting use by unauthorized personnel. The license terms are complex, and the office must manage them carefully to avoid releasing products that infringe the terms.

The geospatial office purchases imagery under the government's framework agreements. Purchasing is conducted through the acquisition desk: the office manager authorizes purchases above the threshold (typically $10,000 per purchase), the technical lead reviews the purchase parameters (the area of interest, the resolution, the spectral bands, the acquisition window, the cloud cover threshold), and the data pipeline operator submits the purchase orders and records the license terms. The license terms are recorded in the catalog and archive submodule and enforced at the production and archive stages, where products are checked against the license terms before release. License enforcement prevents the office from releasing products that infringe the data providers' intellectual property.

The pricing of commercial imagery varies by resolution, by spectral bands, by acquisition type (tasking vs. archive), and by license type. Archive imagery (previously acquired) is typically less expensive than tasked imagery (new acquisitions). Standard license (internal government use) is less expensive than enhanced license (public release permitted). The office manager should be aware of the pricing structure and should select the license type that matches the product's intended use: a product that will be released to the public requires an enhanced license, which is more expensive but permits public release; a product that will be used internally requires only a standard license, which is less expensive but prohibits public release. Selecting the wrong license type either overspends (purchasing an enhanced license for an internal product) or creates a legal risk (purchasing a standard license for a product that will be released publicly). The pricing also varies by area: a 100-square-kilometre tasked image from Maxar WorldView-3 at 0.31m resolution with a standard license may cost $3,500 to $5,000, while the same area from Planet SkySat at 0.5m resolution may cost $1,500 to $2,500. The office manager should select the provider and the resolution that match the application's requirements: an infrastructure change detection application that requires 0.5m resolution may use SkySat rather than WorldView-3, saving $2,000 per acquisition; a structural geological mapping application that requires 0.31m resolution and shortwave infrared bands must use WorldView-3, because SkySat does not provide the required resolution or the shortwave infrared bands. The selection is a technical judgment (what resolution and what bands does the application require?) combined with a financial judgment (what is the least expensive provider that meets the technical requirements?), and the judgment is the office manager's responsibility.

7.4 Sensor Feed Ingestion

Sensor feed ingestion is the process of receiving and processing continuous data streams from ground-based sensors. The capability is distinct from satellite tasking and imagery purchasing because the data is continuous rather than on-demand, and the ingestion is automated rather than manually initiated. Sensor feeds are subscribed to, not tasked, and the data arrives on the sensor's schedule, not the office's.

The sensor feeds that the geospatial office ingests serve four purposes. Rain gauge data provides the rainfall input for the flood risk domain's hydrological models: the rain gauge measures the

precipitation at a specific point, and the model uses the gauge network's data to estimate the rainfall distribution across the basin. River stage data provides the boundary condition for the flood risk domain's hydraulic models: the river stage sensor measures the water level at a specific point along the river, and the model uses the stage data to set the upstream and downstream boundary conditions. Weather station data provides the meteorological input for the food security domain's crop models: the weather station measures the temperature, the humidity, the wind speed, and the solar radiation, and the model uses the weather data to estimate the crop's evapotranspiration and growth. Air quality monitor data provides the environmental input for the energy infrastructure domain's vulnerability assessment: the air quality monitor measures the pollutant concentrations that may affect the infrastructure's corrosion rate.

The United States Geological Survey ingests stream gauge data from the National Streamflow Statistics Program (maintained by the USGS Water Resources Division, which operates approximately 8,500 stream gauges across the United States), weather data from the National Weather Service (which operates approximately 900 Automated Surface Observing System stations), and seismic data from the Advanced National Seismic System (which operates approximately 1,700 seismic stations). The ingestion is automated through standardized data interfaces: the USGS National Water Information System (NWIS) provides stream gauge data through a REST API; the NWS provides weather data through the AWIPS (Advanced Weather Interactive Processing System) data feed; the ANSS provides seismic data through the IRIS (Incorporated Research Institutions for Seismology) data management centre.

The European Flood Awareness System (EFAS), operational since 2005 and managed by the European Commission's Joint Research Centre, ingests river level data from the national hydrological services of the European Union member states, weather data from the national meteorological services (compiled by the European Centre for Medium-Range Weather Forecasts), and snow data from the European Snow and Ice Service. The ingestion is automated through the European Environment Information and Observation Network (Eionet), which provides standardized data interfaces for the exchange of environmental data across the European Union.

The geospatial office subscribes to sensor feeds provided by the government's hydrological, meteorological, and environmental monitoring authorities. The subscription is established through the inter-agency data-sharing framework (treated in Chapter 28), which governs the data exchange between the office and the authorities. The ingestion is automated through the acquisition desk, which polls the sensor feed APIs at configured intervals (typically every 15 minutes for rain gauge and river stage data, every hour for weather station data, every 24 hours for air quality data). The data pipeline operator monitors the ingestion (verifying that the feeds are active and delivering data within the expected time window), the technical lead configures the ingestion parameters (the polling interval, the data format, the quality flags), and the specialists consume the ingested data through the modeling workbench.

Sensor feed ingestion allows the office to produce real-time intelligence, and it distinguishes the office's event-driven production from its scheduled production. The flood risk domain's event-driven production depends on the sensor feeds: when the rain gauge data shows rainfall exceeding a threshold, or the river stage data shows the river rising above a flood level, the event-driven production is activated, and the office produces real-time flood inundation assessments for the

emergency management authority. Without the sensor feeds, the office would be limited to scheduled production (seasonal flood hazard maps, not real-time event response), and the office's value to the emergency management authority would be substantially reduced.

7.5 License Management

License management is the process of recording, tracking, and enforcing the data license terms that govern the use of the data the office acquires. License terms are recorded at the acquisition stage, enforced at the production and archive stages, and audited by the national audit authority. License management is conducted through the catalog and archive submodule: the data pipeline operator records the terms, the production operator enforces them, and the office manager ensures compliance.

License terms vary by data source, and the variation is significant. Free and open data — Landsat imagery (distributed under the NASA Landsat Data Continuity Mission data policy, which places no restrictions on use or redistribution) and Sentinel imagery (distributed under the Copernicus free and open data policy, Regulation (EU) No 377/2014, which requires attribution but permits all uses including commercial) — is licensed under terms that allow free use, including commercial use, with attribution. The office may use free and open data for any purpose, including commercial commissions, without purchasing additional licenses.

Commercial imagery is licensed under terms that are more restrictive and that vary by provider and by license type. Maxar's standard end-user license agreement (EULA) permits internal use by the licensed organization but prohibits redistribution, public display, and the creation of derivative products for public release without an enhanced license. Airbus's EULA has similar restrictions, with additional terms for the Pleiades Neo constellation that distinguish between "Professional Use" (internal analysis and decision-making) and "Public Use" (display or distribution to third parties). Planet's EULA distinguishes between "Analytic" licenses (for internal analysis) and "Display" licenses (for public display in publications and presentations). The license terms may also specify the number of authorized users, the duration of the license (perpetual vs. subscription), and the geographic scope of use.

Sensor feed data is licensed under terms that vary by provider. The USGS stream gauge data is free and open (no restrictions on use). The NWS weather data is free and open. The ANSS seismic data is free and open through the IRIS data management centre. However, sensor feed data from national hydrological and meteorological services in some jurisdictions may be restricted: some services charge for their data and impose license terms that restrict use to the subscribing office, prohibit redistribution, or require attribution. The office must record the license terms for each sensor feed and enforce the terms when the data is used in the office's products.

The license management process has three stages. At the recording stage, the data pipeline operator records the license terms in the catalog and archive submodule at the moment of acquisition. The record includes the data source, the license type, the license restrictions, the license expiration date, and the number of authorized users. At the enforcement stage, the production operator checks each product against the license terms of its source data before the product is released. The check verifies that the product's intended use (public release, internal use, commercial commission) is permitted by the source data's license. If the product's intended use is not permitted, the production operator halts the release and notifies the office manager. At the audit stage, the national audit

authority examines the license records and the enforcement records during the annual audit, and the audit's findings may trigger corrective action where license terms have been violated.

The United States Geological Survey's data management policy, documented in Survey Manual Chapter 502.1, requires that every dataset be recorded with its license terms, enforced through the Survey's publication review process. The Copernicus programme's data policy, documented in Regulation (EU) No 377/2014, requires that every product be checked against the license terms of its source data, enforced through the programme's quality assurance system. License management is the legal integrity guarantee of the pipeline, and it is the guarantee that the national geospatial agencies provide.

7.6 The Provenance Record

The provenance record is the record of the data's source, acquisition, and license. It is the foundation of the provenance chain that runs through the pipeline. The provenance record is established at the acquisition stage and comprises the data source (e.g., Sentinel-2A, acquisition ID S2A_MSIL2A_20240115T060000_N0509_R034_T42SXA_20240115T080000), the acquisition method (subscription, tasking, purchase), the acquisition date (the date the satellite acquired the imagery, not the date the office received it), the ingestion date (the date the office received and ingested the data), the license terms (free and open, commercial standard, commercial enhanced), the intended domain consumer (flood risk, food security, energy infrastructure, minerals), and the acquisition request that initiated the acquisition (the specialist's request or the standing subscription).

The record is maintained in the catalog and archive submodule and is linked to the processing record, which records the processing parameters and the processing date, and to the modeling record, which records the model version and the model run date. The chain of records is the provenance chain, and it allows the office to trace any product back to its source data. The chain is unbroken. Any break in the chain — a missing data source, an unrecorded license term, an undocumented acquisition date — is a defect that must be corrected before the product is released.

The provenance record serves three purposes, each of which has practical consequences for the office. First, it makes the office products auditable. Second, it makes the products reproducible. Third, it makes the products defensible. It makes the office's products auditable: the national audit authority can trace any product back to its source data and verify that the data was acquired legally (under the appropriate license terms), processed correctly (through the standard processing chain with the documented parameters), and modeled appropriately (using the validated model version with the documented model parameters). The audit trail also allows the audit authority to verify that the commercial firewall was maintained: the charge code on the provenance record shows whether the data was acquired for commercial or public-service use, and the audit authority can verify that commercial data was not used in public-service products and vice versa. It makes the office's products reproducible: another analyst can take the provenance record, retrieve the same source data, apply the same processing parameters, run the same model, and reproduce the same product. It makes the office's products defensible: if a consumer challenges the product's accuracy or its legality, the office can produce the provenance record and demonstrate that the product was produced from appropriate data, under appropriate license terms, through appropriate processing and modeling.

The United States Geological Survey's data management policy requires that every dataset be recorded with its source, its acquisition, and its license, enforced through the Survey's data management system. The Copernicus programme's data quality framework requires that every product be traceable to its source data, enforced through the programme's quality assurance system. The provenance record is the integrity guarantee that the national geospatial agencies provide, and it is the guarantee that the pipeline is designed to establish for the geospatial office.

7.7 The Acquisition Discipline

The acquisition discipline is the daily practice of acquiring the right data, at the right time, with the right license, and recording the acquisition in the provenance record. The discipline comprises the daily ingestion check (the data pipeline operator verifies that all overnight acquisitions were received and ingested correctly), the tasking request review (the technical lead reviews the specialist's tasking request for appropriateness before the operator submits it to the provider), the purchase order authorization (the office manager authorizes purchases above the threshold), the license term recording (the operator records the license terms at the moment of acquisition), and the provenance record maintenance (the operator maintains the record through the pipeline's subsequent stages). Each of these practices is specified in the job breakdown sheets in Chapter 15.

The acquisition discipline is maintained through the office's intake function, which determines which acquisition requests are accepted and which are declined. An office that accepts every acquisition request will acquire more data than it can process, and its pipeline will be overloaded. An office that declines too many requests will not acquire enough data, and its pipeline will be underutilized. The intake function is the operational expression of the office's priorities: the office acquires the data that its consumers need for their decisions, not the data that is available or that the office is offered. The intake function is maintained by the office manager with the technical lead configuring and the specialists requesting.

The acquisition discipline is the discipline of the data pipeline operator, who operates the acquisition function on a daily basis. The operator who maintains the discipline acquires the right data (the data the specialists have requested, at the resolution and the timing the applications require), records the acquisition (the data source, the license terms, the acquisition date, and the intended domain consumer), and maintains the provenance record (the unbroken chain from acquisition through archive). The operator's discipline is the foundation of the pipeline's integrity, and it is the discipline that the training pipeline, the qualifications framework, and the professional identity are designed to establish and maintain. The acquisition discipline also includes the management of data gaps. Data gaps occur when a scheduled acquisition is not received (a satellite passes but the image is too cloudy), when a sensor feed goes offline (a rain gauge malfunctions or a communication link fails), or when a commercial provider declines a tasking request (the satellite's orbit does not align with the area of interest within the required window). The operator must identify the gap, assess its impact on the pending analytical work, determine whether alternative data sources are available (a different satellite, a different sensor, a different provider), and communicate the gap and the alternative to the affected specialist. The data gap management is the operational resilience dimension of the acquisition discipline, and it is the dimension that distinguishes an operator who

merely executes the acquisition procedures from an operator who actively manages the pipeline's data supply.

7.8 Worked Example: Acquiring Imagery for a Flood Event Response

At 03:00 on a Wednesday morning, the sensor feed dashboard alerts the data pipeline operator that the river stage at Gauge Station RB-04 on the Riverbourne has risen from 2.1 metres to 3.8 metres in 90 minutes, crossing the 5-year flood threshold of 3.5 metres. The rain gauge network shows 120 millimetres of rainfall in the preceding 6 hours over the upper basin. The operator notifies the technical lead, who notifies the office manager. The office manager activates the incident response procedure (section 13.6) and assigns the hazard specialist to produce a flood inundation assessment.

The hazard specialist identifies the data requirements for the assessment. The specialist needs: (1) the latest sensor feed data (rain gauge and river stage, already being ingested continuously); (2) the most recent satellite imagery of the basin (to map the current flood extent from space); (3) the LiDAR-derived digital elevation model (already in the spatial data infrastructure, acquired in 2019); (4) the cadastre feed (already in the spatial data infrastructure, synchronized weekly). The specialist's data requirements for the event are limited to the satellite imagery, because the sensor feeds and the base layers are already available.

The specialist checks the recent satellite acquisitions. A Sentinel-2 image was acquired over the basin at 06:00 UTC on the same morning, approximately 3 hours after the flood threshold was crossed. The image has 70 per cent cloud cover (the storm that caused the flooding is still over the basin), which makes the optical imagery unsuitable for flood extent mapping. The specialist checks for Sentinel-1 (synthetic aperture radar) imagery, which can penetrate clouds and map the flood extent regardless of cloud cover. A Sentinel-1 image was acquired over the basin at 05:30 UTC (Sentinel-1's C-band SAR can detect standing water through clouds, because calm water reflects the radar pulse away from the sensor and appears dark in the image). The Sentinel-1 image is available through the Copernicus Open Access Hub and can be downloaded free of charge under the Copernicus free and open data policy.

The specialist submits a request for the Sentinel-1 image through the acquisition desk. The data pipeline operator retrieves the image from the Copernicus Open Access Hub (the image was acquired at 05:30 UTC and made available on the Hub at 07:45 UTC, approximately 2 hours and 15 minutes after acquisition). The operator ingests the image into the processing engine under the public-service charge code (PUB-2024-119), records the license terms (Copernicus free and open data, commercial use permitted with attribution), and records the provenance: data source (Sentinel-1B, acquisition ID S1B_IW_SLC__1SDV_20240117T053000_20240117T053030_034182_041A5B_9C8A), acquisition method (subscription, free and open), acquisition date (2024-01-17 05:30 UTC), ingestion date (2024-01-17 08:15 UTC), license terms (free and open, attribution required), intended domain consumer (flood risk).

The specialist also identifies the need for high-resolution optical imagery for the post-event damage assessment, which will be conducted after the flood waters recede. The specialist submits a tasking request through the acquisition desk for a Maxar WorldView-3 image of the inundated area, with 0.31-metre panchromatic and 1.24-metre multispectral resolution, 8-band multispectral plus 8-band shortwave infrared, acquisition window 24-48 hours after the flood waters recede, maximum 10

per cent cloud cover. The technical lead reviews the tasking parameters and confirms that they are appropriate for a post-event damage assessment. The office manager authorizes the purchase, which is estimated at $4,200 under the government's framework agreement with Maxar (archive imagery, standard license, internal government use).

The data pipeline operator submits the tasking request to Maxar through the acquisition desk. Maxar confirms the tasking and provides an expected acquisition window of 48-72 hours, subject to cloud cover. The operator records the tasking request, the expected acquisition window, the price ($4,200), and the license terms (Maxar standard EULA, internal government use, no redistribution, no public display without enhanced license) in the catalog and archive submodule.

The specialist also requests a Planet SkySat image for the interim period (between the flood's peak and the Maxar acquisition), to provide a preliminary damage assessment. The SkySat image (0.5m panchromatic, 2m multispectral) is less detailed than the WorldView-3 but is available more quickly (SkySat's constellation of 21 satellites provides a 1-2 day revisit capability). The office manager authorizes the purchase ($1,800, standard license, internal government use), and the operator submits the tasking request to Planet.

The acquisition process for the flood event response illustrates the acquisition function's four sub-capabilities operating together. Satellite tasking is used for the Maxar and Planet imagery (new acquisitions, tasked for specific applications). Imagery purchasing is used for the Maxar and Planet imagery (commercial imagery, purchased under framework agreements). Sensor feed ingestion is used for the rain gauge and river stage data (continuous data streams, already being ingested). License management is used for all data (the Sentinel-1 image is free and open; the Maxar and Planet imagery is commercial, with standard licenses for internal government use). The provenance record is established for all data at the moment of acquisition, and the record will be maintained through the processing, modeling, production, and archive stages as the flood inundation assessment is produced.

Chapter 8 — The Spatial Baseline

8.1 The Base Layers

The spatial baseline is the geospatial office's maintained set of base layers: the digital elevation model, land cover, hydrography, administrative boundaries, and the cadastre feed. The base layers provide the geographic context for every analysis the office conducts. Their integrity is the non-negotiable precondition for every product. The base layers are maintained by the data pipeline operator, with the technical lead configuring the maintenance parameters and the specialists consuming the base layers through the modeling workbench.

The digital elevation model (DEM) is the three-dimensional representation of the Earth's surface. The flood risk domain's hydrological modeling depends on it: the DEM provides the terrain surface that the hydraulic model uses to simulate the flow of water across the floodplain. The DEM's resolution determines the hydraulic model's accuracy: a 1-metre LiDAR DEM can resolve individual levees, road embankments, and building footprints that affect the flood flow, while a 30-metre SRTM DEM smooths over these features and produces a less accurate flood simulation. The DEM's vertical accuracy is equally important: a vertical error of 1 metre in the DEM translates to a horizontal error of tens of metres in the flood inundation extent, because the floodplain is typically very flat and a small vertical error moves the water's edge a large horizontal distance.

The land cover layer classifies the Earth's surface by land cover type: forest, grassland, cropland, wetland, urban, water, bare soil, snow and ice. The food security domain's crop stratification depends on it: the land cover layer identifies which areas are cropland, and the crop condition assessment is stratified by the cropland parcels. The energy infrastructure domain's encroachment detection depends on it: the land cover layer identifies the vegetation type near the infrastructure, and the encroachment assessment uses the vegetation type to estimate the growth rate and the encroachment risk. The flood risk domain's roughness estimation depends on it: the land cover layer identifies the surface roughness (Manning's n values), which the hydraulic model uses to estimate the flow resistance. Urban areas have high roughness (n = 0.035-0.060); forests have moderate roughness (n = 0.080-0.120); cropland has low roughness (n = 0.025-0.035); water surfaces have minimal roughness (n = 0.020-0.030).

The hydrography layer represents the Earth's surface waters: rivers, streams, lakes, reservoirs, canals, and coastlines. The flood risk domain's flow modeling depends on it: the hydrography layer provides the channel network that the hydraulic model uses to route the flow. The food security domain's irrigation assessment depends on it: the hydrography layer identifies the water sources that the crops depend on. The minerals domain's environmental assessment depends on it: the hydrography layer identifies the watercourses that may be affected by mining activities. The hydrography layer includes the channel geometry (the cross-sections, the longitudinal profiles), the flow directions, the stream orders (the Strahler stream order, which classifies streams from first-order headwater streams to the highest-order main stem), and the naming (the official names of the water features, which are used in the office's products for consumer orientation).

The administrative boundaries layer represents the jurisdictional boundaries: national, provincial, municipal, and other administrative divisions. All domains' reporting and aggregation depend on it: the flood exposure report aggregates the affected population by municipality; the crop

condition bulletin reports the yield forecast by province; the infrastructure monitoring report identifies the infrastructure by jurisdiction. The administrative boundaries layer is maintained by the national mapping authority and updated as jurisdictions are reorganized (through municipal mergers, provincial boundary adjustments, or the creation of new administrative units).

The cadastre feed represents the land parcels and their legal attributes: the parcel identifier, the parcel boundary, the owner (or the ownership type: private, state, communal), the land use designation, the encumbrances (easements, mortgages, liens), and the transaction history. All domains' parcel-level analysis depends on it: the flood exposure report identifies the affected parcels and their owners; the crop condition assessment stratifies the vegetation indices by parcel; the infrastructure monitoring report identifies the parcels along the infrastructure corridors; the mineral prospectivity study identifies the parcels with exploration licenses. The cadastre feed connects the office's intelligence to the legal and administrative reality of the land.

8.2 The Digital Elevation Model in Detail

The digital elevation model deserves detailed treatment because it is the base layer that most directly determines the quality of the flood risk domain's products, and because the choice of DEM source involves trade-offs between resolution, coverage, cost, and currency that the office must understand.

Global DEMs provide worldwide coverage at moderate resolution. The Shuttle Radar Topography Mission (SRTM) DEM, acquired in February 2000 during an 11-day Space Shuttle mission (STS-99), provides near-global coverage (56°S to 60°N) at 1 arc-second (approximately 30 metres) resolution. SRTM was a single-pass interferometric synthetic aperture radar (InSAR) mission: two radar antennas, one in the Shuttle's payload bay and one on a 60-metre mast, acquired the same terrain from slightly different angles, and the phase difference between the two radar returns was used to calculate the elevation. SRTM's vertical accuracy is approximately 5-10 metres (90 per cent linear error), which is sufficient for regional flood modelling but insufficient for urban flood modelling where road embankments and levees must be resolved. The SRTM DEM is freely available through the USGS EarthExplorer.

The Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) Global Digital Elevation Model (GDEM), released in 2009 and updated in 2011 (Version 2) and 2019 (Version 3), provides global coverage at 1 arc-second (30 metres) resolution. ASTER GDEM is produced from stereoscopic optical imagery acquired by the ASTER instrument on NASA's Terra satellite, launched in 1999. ASTER GDEM's vertical accuracy is similar to SRTM's (approximately 7-14 metres), but ASTER GDEM covers the latitudes beyond SRTM's coverage (up to 83°N and 83°S). ASTER GDEM is freely available through the NASA Earthdata portal.

The Copernicus DEM (GLO-30), released in 2020 and derived from the WorldDEM dataset (acquired by the Airbus TanDEM-X mission, 2010-2015), provides global coverage at 1 arc-second (30 metres) resolution with improved vertical accuracy (approximately 2.5 metres root mean square error, a significant improvement over SRTM and ASTER). The Copernicus DEM is freely available through the Copernicus Open Access Hub and has replaced SRTM and ASTER as the preferred global DEM for many applications.

National high-resolution DEMs, typically derived from airborne LiDAR (Light Detection and Ranging), provide much higher resolution and accuracy. LiDAR acquires elevation data by measuring the time of flight of laser pulses from an aircraft to the ground and back. A typical LiDAR acquisition produces a point cloud with 2-8 points per square metre, which is processed into a DEM at 1-metre or sub-metre resolution with vertical accuracy of 0.10-0.15 metres (15 centimetres root mean square error). LiDAR DEMs can resolve individual levees, road embankments, building footprints, and drainage ditches that affect flood flow, and they are the preferred DEM for urban flood modelling. The United States' 3D Elevation Program (3DEP), administered by the USGS, is acquiring LiDAR coverage for the entire United States at 1-metre resolution (Quality Level 2, 0.10m RMSE), with a target completion date of 2023 (later revised). The United Kingdom's Environment Agency maintains a LiDAR DEM for England at 1-metre and 50-centimetre resolution, freely available through the Survey Open Data portal. The Netherlands' AHN (Actueel Hoogtebestand Nederland) provides nationwide LiDAR coverage at 0.5-metre resolution, updated every 3-5 years.

The office's DEM selection depends on the application. For regional flood modelling (basin-scale, 100-square-kilometre or larger), the Copernicus DEM at 30 metres is sufficient. For urban flood modelling (city-scale, where road embankments and levees must be resolved), LiDAR at 1 metre or better is required. For crop condition assessment, the DEM is used primarily for terrain correction (removing the illumination effects caused by terrain slope and aspect), and the Copernicus DEM at 30 metres is sufficient. For infrastructure monitoring, the DEM is used for line-of-sight analysis and for slope stability assessment, and LiDAR at 1 metre or better is preferred. The office should maintain both a global DEM (for regional applications) and a national LiDAR DEM (for urban and infrastructure applications), and the specialists should select the appropriate DEM for each application. The DEM may also be available in different forms: a digital surface model (DSM), which includes all features on the surface (buildings, trees, infrastructure), and a digital terrain model (DTM), which represents the bare earth surface (with buildings and trees removed). LiDAR data can produce both a DSM and a DTM from the same point cloud: the DSM is derived from the first returns (the first surface the laser pulse hits), and the DTM is derived from the last returns (the last surface the laser pulse hits, which is typically the ground). Flood modelling requires the DTM, because the flood water flows on the ground surface, not on the rooftops. Infrastructure line-of-sight analysis requires the DSM, because the line of sight may be blocked by buildings and trees. The office must maintain both forms and select the appropriate form for each application.

8.3 The Cadastre Feed

The cadastre feed is the feed from the cadastral authority that provides the geospatial office with the current state of the cadastre — the official, legally binding map of land boundaries and titles. The cadastre feed makes the office's products actionable. A flood inundation map is meaningful when it can be associated with specific parcels. A crop condition assessment is meaningful when it can be associated with specific fields. An infrastructure risk assessment is meaningful when it can be associated with specific properties. A mineral prospectivity study is meaningful when it can be associated with specific exploration licenses. The cadastre feed connects the office's intelligence to the legal and administrative reality of the land.

The cadastre feed is maintained by the cadastral authority, not by the geospatial office. The office consumes the feed as it is provided. The office does not modify the cadastre, does not determine land boundaries, and does not resolve boundary disputes. The office uses the cadastre as the legally authoritative record of land boundaries and titles, and it reports any discrepancies to the cadastral authority through the formal discrepancy reporting channel (treated in Chapter 23). The distinction between the cadastral authority and the geospatial office is treated in Chapter 3, and the interface between the two is treated in Chapter 23.

The cadastre feed's data model typically includes the parcel identifier (a unique identifier, such as the cadastral parcel number or the National Parcel ID), the parcel geometry (the polygon representing the parcel's boundary, in a standard coordinate reference system), the parcel attributes (the area, the land use designation, the ownership type), and the transaction history (the dates and types of transactions that have affected the parcel: subdivisions, consolidations, transfers, easements). The data model is specified by the cadastral authority and is standardized through the national spatial data infrastructure's data standards.

The French Institut national de l'information géographique et forestière maintains the cadastre (the Plan Cadastral Informatisé) and provides the cadastre feed through the IGN's Géoservices API, which provides Web Feature Service (WFS) and Web Map Service (WMS) interfaces for the cadastre data. The cadastre is updated continuously as land transactions are registered with the French land registry (the Service de la Publicité Foncière). The Singapore Land Authority maintains the Singapore Cadastral Survey Plan and provides the cadastre feed through the SLA's Integrated Land Information Service (INLIS), which provides the parcel data through a web portal and an API. The Netherlands' Cadastre, Land Registry and Mapping Agency (Kadaster) maintains the cadastre (the Kadastrale Kaart) and provides the cadastre feed through the PDOK (Publieke Dienstverlening Op de Kaart) platform, which provides WFS and WMS interfaces for the cadastre and other base layers. The United States does not have a national cadastre; cadastral records are maintained at the county level, and the Bureau of Land Management's National Integrated Land System (NILS) provides a partial federal cadastre for public lands.

The cadastre feed is consumed by the geospatial office through a standardized data interface, typically a Web Feature Service (WFS) or a file download. The WFS interface allows the office to query the cadastre by area (retrieving all parcels within a bounding box), by attribute (retrieving a specific parcel by its identifier), or by spatial filter (retrieving all parcels that intersect a flood inundation polygon). The WFS interface returns the data in a standard format, typically Geography Markup Language (GML), which the office ingests into its spatial data infrastructure. The file download provides the complete cadastre as a bulk download, typically in shapefile or GeoPackage format, which the office ingests periodically (daily or weekly) to maintain a local copy of the cadastre for offline use and for performance reasons (querying a local copy is faster than querying a remote WFS). The cadastre feed's update frequency varies by jurisdiction. In jurisdictions with continuous cadastral updating (France, Singapore, Netherlands), the feed is updated daily or weekly as transactions are registered. In jurisdictions with periodic cadastral updating (some jurisdictions update the cadastre only every few years), the feed is updated less frequently, and the office may be working with a cadastre that is months or years out of date. The office must be aware of the feed's

update frequency and must verify, through the synchronization check (section 8.5), that the cadastre feed used in the production is the most recent version available.

8.4 Base Layer Integrity

Base layer integrity is the principle that no product enters production against an unsynchronized base layer. The principle is maintained through the synchronization check, conducted before every production run, which verifies that the base layers used in the production are synchronized with the current state of the cadastre and the other base layer sources.

The base layers change over time, and products produced against outdated base layers are wrong. The cadastre changes as parcels are subdivided (a 50-hectare agricultural parcel is subdivided into five 10-hectare residential parcels), consolidated (three adjacent parcels are merged into one for a commercial development), or redefined (a boundary dispute is resolved and the boundary is adjusted). The land cover changes as land use changes (a forested area is cleared for agriculture, an agricultural area is developed for housing). The hydrography changes as watercourses are modified (a river is channelized, a canal is constructed, a reservoir is created). The administrative boundaries change as jurisdictions are reorganized (municipalities merge, provinces are redrawn, new administrative units are created). The DEM changes as new LiDAR acquisitions are made (a post-flood LiDAR survey may reveal changes to the channel geometry caused by the flood).

Products produced against outdated base layers are wrong in ways that are not always visible to the consumer. A flood inundation map produced against an outdated cadastre may show a parcel as agricultural when it has been subdivided into residential plots, and the exposure report will undercount the affected population. A crop condition assessment produced against an outdated land cover layer may include non-cropland areas (which have been converted to urban use since the land cover was last updated) and will produce an inaccurate yield forecast. An infrastructure monitoring report produced against an outdated administrative boundaries layer may attribute an encroachment to the wrong jurisdiction, and the report will be sent to the wrong authority. Misleading intelligence — intelligence that is wrong but trusted — is worse than no intelligence, because it is acted upon, and the actions taken on the basis of wrong intelligence may be harmful.

The United States Geological Survey's National Map programme maintains a continuous update process for its base layers. The 3D Elevation Program acquires new LiDAR data on a rolling basis, with each state's LiDAR updated approximately every 8 years. The National Hydrography Dataset is updated through stewardship agreements with state and local agencies that report changes to the hydrography. The National Land Cover Database is updated approximately every 5 years (the most recent version is NLCD 2021, released in 2023). The Survey's analytical products are produced against the current state of the National Map. The Copernicus programme's land monitoring service maintains the CORINE Land Cover (updated every 6 years: 1990, 2000, 2006, 2012, 2018) and the Copernicus Global Land Cover (updated annually, at 100-metre resolution). Geoscience Australia maintains the 1-second SRTM-derived DEM and the national geological map, with updates as new data becomes available.

8.5 The Synchronization Check

The synchronization check is the procedure by which the production operator verifies that the base layers used in a production run are synchronized with the current state of the base layer sources. The check is conducted before every production run, without exception. It comprises the comparison of the base layer version used in the production with the current base layer version maintained by the source. The comparison is automated through the catalog and archive submodule, which maintains the version history of every base layer and flags any production run that uses an outdated version.

The production operator conducts the check, with the technical lead configuring the check parameters and the office manager overseeing the check's execution. The operator retrieves the version number and the version date of each base layer used in the production from the spatial data infrastructure submodule. The operator queries each source authority's data interface for the current version number and version date. If the versions match, the base layer is synchronized and the production may proceed. If the versions do not match, the base layer is outdated, and the operator halts the production and notifies the technical lead.

The technical lead investigates the mismatch. The mismatch may be a synchronization failure (the office's version is outdated because the update was not ingested) or a legitimate version difference (the source authority has released a new version that the office has not yet ingested, but the version the office is using is still valid for the production). If the mismatch is a synchronization failure, the technical lead directs the data pipeline operator to download and ingest the current version, and the production is rescheduled. If the mismatch is a legitimate version difference, the technical lead may authorize the production to proceed with the current version, with the version difference documented in the provenance record.

The synchronization check must not be waived, regardless of the pressure to release a product quickly. The pressure to waive the check will arise, particularly during flood events when the demand for intelligence is urgent and the emergency management authority is requesting the inundation map within hours. An office that waives the check during a flood event produces intelligence based on base layers that may be outdated, and the intelligence may be wrong in ways that are not visible to the consumer. An office that maintains the check during a flood event produces intelligence based on the current base layers. The intelligence will be right, even if it is produced a few minutes later than the waiving would have allowed. The few minutes are the cost of the integrity, and the cost is worth paying.

The office manager oversees the synchronization check and ensures that it is conducted for every production run. The manager reviews the check's results as part of the morning production review (section 13.3) and as part of the product review before release (section 16.3). The manager's oversight ensures that the check is not skipped under pressure and that any mismatches are resolved before the product is released. The synchronization check is documented in the catalog and archive submodule as part of the provenance record. The documentation includes the base layer versions checked, the source authority versions compared, the check result (synchronized or mismatch), and any resolution applied (update ingested or legitimate version difference authorized). The documentation is the audit trail that the national audit authority can examine to verify that the base layer integrity principle was maintained for every production run.

8.6 The National Spatial Data Infrastructure

The national spatial data infrastructure (NSDI) is the national framework for the production, sharing, and governance of spatial data. It comprises the base layer sources, the data sharing agreements, the metadata standards, and the governance arrangements that coordinate the production and sharing of spatial data across the government. The geospatial office's spatial baseline is a participant in the framework, not a parallel structure. The office's base layers are maintained through the framework's sources, agreements, standards, and governance.

The United States Federal Geographic Data Committee (FGDC), established in 1990 by the Office of Management and Budget Circular A-16, coordinates the National Spatial Data Infrastructure. The NSDI was formally established by Executive Order 12906 in 1994, which directed federal agencies to document their spatial data using the FGDC Content Standard for Digital Geospatial Metadata, to make their data available through a National Geospatial Data Clearinghouse, and to participate in the development of the National Spatial Data Infrastructure. The FGDC is composed of representatives from 32 federal agencies and is chaired by the Secretary of the Interior. The NSDI's base layer sources include the USGS (elevation, hydrography, land cover, orthoimagery), the Bureau of the Census (governmental unit boundaries), the Bureau of Land Management (cadastral), the National Oceanic and Atmospheric Administration (geodetic control, bathymetry), and the Natural Resources Conservation Service (soils).

The European Union's INSPIRE Directive (Directive 2007/2/EC), adopted in 2007 and transposed into national law by the member states by 2009, establishes the Infrastructure for Spatial Information in the European Community. INSPIRE defines 34 spatial data themes (annexes I, II, and III), including coordinate reference systems, geographical grid systems, geographical names, administrative units, addresses, cadastral parcels, transport networks, hydrography, protected sites, elevation, land cover, orthoimagery, geology, land use, human health and safety, and environmental monitoring facilities. Each theme has a data specification that defines the data model, the metadata requirements, and the network service requirements. The INSPIRE geoportal provides a single point of access to the spatial data of all European Union member states.

The Australian Spatial Data Infrastructure (ASDI), coordinated by the Intergovernmental Committee on Surveying and Mapping (ICSM), provides the national framework for Australia. The ASDI's base layer sources include Geoscience Australia (elevation, geology, satellite imagery), the Bureau of Meteorology (weather and climate), the Australian Bureau of Statistics (statistical boundaries), and the state land agencies (cadastral). The Singapore Geospatial Collaborative Environment (SG-SPACE), coordinated by the Singapore Land Authority, provides the national framework for Singapore. The model is consistent across comparative practice: the national spatial data infrastructure provides the framework, and the government's analytical agencies participate in it.

The geospatial office's participation in the NSDI is the mechanism by which the office's base layers are maintained at the current state of the national spatial record. The office subscribes to the base layer sources (through the data sharing agreements), follows the data sharing agreements (which specify the terms of use, the attribution requirements, and the update notifications), complies with the metadata standards (which ensure that the office's products can be discovered and used by other agencies), and participates in the governance arrangements (which coordinate the production and sharing of spatial data across the government). An office that does not participate maintains its own

base layers in isolation. Isolated base layers will diverge from the national spatial record, and the divergence will produce intelligence that is inconsistent with the national spatial record.

8.7 The Governance of the Base Layers

The governance of the base layers is the arrangement by which the base layers are maintained, updated, and quality-assured. It comprises the base layer sources, the update cycles, the quality assurance processes, and the dispute resolution mechanisms.

The base layer sources are the institutions that maintain the base layers. The national mapping agency maintains the DEM, the land cover, the orthoimagery, and the administrative boundaries. The cadastral authority maintains the cadastre. The national hydrological service maintains the hydrography (in some jurisdictions; in others, the hydrography is maintained by the mapping agency). The national geological survey maintains the geological map (which is not a base layer in the strict sense but is a thematic layer used by the minerals domain).

The update cycles vary by base layer. The cadastre is updated continuously as land transactions occur (in jurisdictions with continuous cadastral updating) or periodically (in jurisdictions with periodic cadastral updating). The DEM is updated as new LiDAR acquisitions are made, typically every 5-10 years for national programmes. The land cover is updated every 3-6 years (CORINE) or annually (Copernicus Global Land Cover). The hydrography is updated as changes are detected (through field surveys, satellite imagery analysis, or user reports). The administrative boundaries are updated as jurisdictions are reorganized, which occurs irregularly and typically requires legislative action.

The quality assurance processes are the processes by which the base layers' accuracy is verified. Ground truth validation compares the base layer against field measurements: a DEM may be validated against GPS survey points; a land cover layer may be validated against field plots; a hydrography layer may be validated against field surveys of the channel geometry. Cross-referencing with other sources compares the base layer against independent datasets: a DEM may be cross-referenced against a different DEM (e.g., SRTM vs. Copernicus DEM); a land cover layer may be cross-referenced against satellite imagery. User feedback mechanisms allow the users of the base layers to report errors: the USGS National Map Corps allows volunteers to report errors in the National Map's structural and hydrographic features, and the reports are verified and incorporated into the next update.

The geospatial office participates in the governance by reporting discrepancies (the office's analytical work may reveal base layer errors that the source authority's own processes have not identified), by participating in the quality assurance processes (the office may contribute ground truth data from its field work), and by providing feedback on the base layers' fitness for the office's purposes (the office may report that the DEM's resolution is insufficient for urban flood modelling, or that the land cover's classification scheme does not distinguish between crop types). The office benefits by receiving the current base layers, by participating in the governance arrangements, and by having a voice in the base layers' development. The dispute resolution mechanisms are the mechanisms by which disagreements about the base layers are resolved. Disputes may arise when two source authorities provide conflicting data for the same feature, when a user reports an error that the source authority disputes, or when the update cycle of one base layer is incompatible with the

update cycle of another. The dispute resolution mechanisms are typically coordinated through the national spatial data infrastructure governance arrangements, which provide a forum for the source authorities and the users to discuss and resolve the disputes. The base layer governance is not a one-time setup but a continuous relationship between the geospatial office and the source authorities, maintained through the inter-agency data-sharing framework, the weekly base layer integrity audit, and the discrepancy reporting channel. The relationship has a contractual dimension (the data sharing agreement specifies the terms of use and the update notifications), a technical dimension (the data interfaces through which the base layers are exchanged), and an institutional dimension (the governance arrangements through which the parties coordinate and resolve disputes). All three dimensions must be maintained for the spatial baseline to function.

8.8 Worked Example: Resolving a Cadastre Discrepancy

During the production of a flood exposure assessment for the Riverbourne Basin, the hazard specialist overlays the flood inundation map with the cadastre feed and identifies 347 affected parcels. The specialist notices that Parcel RB-2019-0476, a 12-hectare parcel on the eastern bank of the Riverbourne, appears as a single agricultural parcel in the cadastre, but the satellite imagery (Sentinel-2, acquired the previous week) shows the parcel divided into two areas: a 8-hectare area that is still agricultural (cropland, with visible crop rows) and a 4-hectare area that has been developed (with a newly constructed building visible in the imagery). The cadastre shows the parcel as agricultural; the imagery shows part of it as developed.

The specialist reports the discrepancy to the office manager, who records the discrepancy in the office's discrepancy log and directs the specialist to submit a formal discrepancy report to the cadastral authority. The discrepancy report is submitted through the formal channel (the inter-agency data-sharing framework's discrepancy reporting procedure, treated in Chapter 23). The report includes: the parcel identifier (RB-2019-0476), the parcel's location (eastern bank of the Riverbourne, coordinates 42.351°N, 7.824°E), the cadastre's record (single agricultural parcel, 12 hectares), the observed condition (4-hectare area developed with a new building, visible in Sentinel-2 imagery acquired 2024-01-15), and the supporting evidence (the Sentinel-2 image, annotated to show the developed area).

The cadastral authority receives the discrepancy report and initiates an investigation. The investigation reveals that the parcel was subdivided in October 2023 (three months before the discrepancy was reported): the 4-hectare area was sold to a commercial developer, and a building permit was issued for the construction. The subdivision was registered with the land registry, but the cadastre had not been updated to reflect the subdivision because of a backlog in the cadastral updating process.

The cadastral authority resolves the discrepancy by updating the cadastre: the original parcel (RB-2019-0476) is subdivided into two parcels: RB-2019-0476-A (8 hectares, agricultural) and RB-2019-0476-B (4 hectares, commercial, with a building). The updated cadastre is released through the cadastre feed, and the geospatial office ingests the update during the next synchronization check.

The office's flood exposure assessment is updated to reflect the corrected cadastre. The original assessment identified 347 affected parcels. The corrected assessment identifies 348 affected parcels (the original parcel has been split into two), and the exposure report is updated: the 4-hectare

commercial parcel (RB-2019-0476-B) has a building that was not identified in the original assessment, and the building's occupants are added to the affected population estimate. The corrected assessment is delivered to the emergency management authority with a note explaining the correction and its basis.

This worked example illustrates the spatial baseline's two-way relationship with the cadastral authority. The cadastral authority provides the cadastre, and the geospatial office provides the discrepancy reporting that helps the cadastre stay current. The office did not modify the cadastre (the modification was the cadastral authority's function), and the office did not produce the exposure assessment against the corrected cadastre before the correction was made (the office waited for the cadastral authority's resolution before adjusting its products). The boundary between the cadastral function and the analytical function was maintained, and the discrepancy reporting was the mechanism by which the office contributed to the cadastre's quality without compromising the cadastre's legal authority.

Chapter 9 — Processing

9.1 The Processing Function

Processing is the third capability of the pipeline. It transforms raw acquired data into analysis-ready products through georeferencing, calibration, mosaicking, projection management, cloud masking, and data fusion. The capability is shared across all four domain applications, configured by the technical lead, and operated by the data pipeline operator. Every processed product is validated against independent ground control points before it is released to the modeling workbench. Geometric errors introduced at the processing stage propagate through modeling and production, and the consumer sees them as products that point at the wrong place on the ground.

The processing function is conducted through the processing engine submodule of the geospatial information system, treated in Chapter 29. The engine manages the processing queue, the parameter library, the ground control point library, and the output catalog. The data pipeline operator configures the processing runs, the technical lead validates the outputs against independent check points, and the specialists consume the processed data through the modeling workbench. The design and use of the processing engine are specified in the job breakdown sheets in Chapter 15.

The processing cycle operates on a daily rhythm. The data pipeline operator reviews the processing queue at the start of the shift, prioritizes the runs that the modeling workbench requires for the day's production, and configures the parameters for each run. The technical lead reviews the validation results for each completed run before the product is released, and any product that fails the validation is reprocessed with corrected parameters or rejected. The cycle ensures the modeling workbench always has a queue of validated, analysis-ready data, and that no unvalidated data enters the modeling stage.

The processing capability has six sub-capabilities, each treated in a separate section of this chapter. Georeferencing (section 9.2) associates raw imagery with geographic coordinates. Calibration (section 9.3) adjusts the radiometric values to a standard reference. Mosaicking (section 9.3) combines multiple overlapping images into a single seamless image. Projection management (section 9.4) maintains the coordinate reference systems and transforms data between projections as required. Cloud masking (section 9.5) identifies and removes clouds and cloud shadows from optical imagery. Data fusion (section 9.6) combines data from multiple sources into a single integrated product. The geometric integrity guarantee (section 9.7) validates the processed products against independent ground control points.

9.2 Georeferencing

Georeferencing associates raw imagery with geographic coordinates, transforming an image from a pixel-based representation to a spatially referenced one. The process is the foundation of every subsequent processing step: mosaicking, projection management, cloud masking, and data fusion all require georeferenced imagery. The process identifies ground control points (GCPs) in the imagery, matches them to known geographic coordinates, and applies a transformation model that maps pixel coordinates to geographic coordinates.

Ground control points are features in the imagery whose geographic coordinates are known independently. A GCP may be a road intersection, a building corner, a bridge, or any other feature that is visible in the imagery and whose coordinates have been determined by a higher-accuracy method (GNSS survey, surveyed control point, or reference imagery of known accuracy). The GCPs must be well distributed across the image: clustered GCPs produce a transformation that is accurate in the cluster area but inaccurate elsewhere, because the transformation is over-constrained in one area and under-constrained in another. A minimum of four GCPs is required for a first-order polynomial, seven for a second-order, and ten for a third-order, but the practical minimum is higher (typically 6-8 for first-order, 12-15 for second-order) to provide redundancy and to allow the identification of outliers.

The transformation model is typically a polynomial of first, second, or third order, or a spline. A first-order polynomial (affine transformation) corrects for translation (shift in X and Y), rotation, and scale: six parameters, requiring a minimum of three GCPs. A second-order polynomial corrects for additional second-order distortions (such as perspective and non-uniform scale): twelve parameters, requiring a minimum of six GCPs. A third-order polynomial corrects for higher-order distortions: twenty parameters, requiring a minimum of ten GCPs. A spline (such as a thin-plate spline) corrects for local distortions that the polynomial models cannot represent, by fitting the transformation exactly at each GCP and interpolating smoothly between them. The spline is useful when the distortion varies across the image (for example, in mountainous terrain where the terrain-induced parallax varies with the slope and aspect), but it is more computationally intensive and can overfit the GCPs if they contain errors.

The choice of transformation model depends on the distortion characteristics. Satellite imagery that has been geometrically corrected by the data provider (such as Sentinel-2 Level-1C or Landsat Collection 2 Level-1) typically requires only a first-order polynomial or no additional georeferencing, because the provider has already applied a rigorous orthorectification using the satellite's orbital model and a DEM. Aerial photography that has not been orthorectified may require a second-order or third-order polynomial, because the perspective distortion varies across the image. Historical imagery that was scanned from film may require a spline, because the film distortion is non-uniform and cannot be modeled by a polynomial.

The accuracy of the georeferencing is determined by the number, distribution, and accuracy of the GCPs, and by the appropriateness of the transformation model. The accuracy is measured by the root mean square error (RMSE) at the GCPs and, more importantly, at independent check points (ICPs) that are not used in the transformation. The RMSE at the GCPs measures how well the transformation fits the control points; the RMSE at the ICPs measures how well the transformation generalizes to the rest of the image. A transformation with a low GCP RMSE but a high ICP RMSE is overfit: it passes through the control points but is inaccurate elsewhere. The ICP RMSE is the measure that the technical lead uses to validate the georeferencing, and it is the measure that the processing standards specify.

The United States Geological Survey's Earth Resources Observation and Science Center georeferences Landsat imagery using a global ground control point library maintained by the Survey. The georeferencing accuracy is specified in the Landsat Collection 2 Level-1 product: the RMSE is required to be less than 12 metres against the ground control point library. The European Space

Agency's Copernicus programme georeferences Sentinel-2 imagery against a global ground control point library maintained by the programme; the Level-1C product requires a sub-pixel georeferencing accuracy (less than 10.3 metres RMSE at the 95 per cent confidence level, which is less than one pixel of the Sentinel-2 10-metre bands). The geospatial office maintains its own ground control point library through the catalog and archive submodule, drawing on the cadastre, on surveyed control points provided by the national mapping agency, and on high-precision reference imagery.

9.3 Calibration and Mosaicking

Calibration adjusts the radiometric values of the imagery to a standard reference, transforming the raw digital numbers (DNs) into physically meaningful values such as reflectance or radiance. The calibration is conducted using the sensor's calibration coefficients, which are provided by the satellite operator and which account for the sensor's radiometric response, the solar irradiance, and the atmospheric conditions.

The calibration proceeds in two stages. Top-of-atmosphere (TOA) reflectance corrects for the sensor response and the solar irradiance, converting the DNs to a value that represents the proportion of incoming solar radiation reflected by the Earth's surface and the atmosphere. The TOA reflectance is computed from the DNs using the formula: ρ_TOA = (M_L × DN + A_L) / (sin(θ_SE)), where M_L is the band-specific multiplicative rescaling factor, A_L is the band-specific additive rescaling factor, and θ_SE is the sun elevation angle. The coefficients are provided in the imagery's metadata (for Landsat, in the MTL file; for Sentinel-2, in the MTD_MSIL1C.xml file).

Surface reflectance (SR) additionally corrects for atmospheric scattering and absorption, converting the TOA reflectance to a value that represents the proportion of solar radiation reflected by the Earth's surface alone. The atmospheric correction requires auxiliary atmospheric data: aerosol optical thickness (the amount of particulate matter in the atmosphere), water vapor (the amount of atmospheric water), and ozone (the amount of atmospheric ozone). These data are obtained from meteorological sources (such as the MODIS atmospheric products, which provide aerosol optical thickness at 10-kilometre resolution) or estimated from the imagery itself (through algorithms such as the Dark Object Subtraction method, which estimates the atmospheric haze from the darkest pixel in the image). The atmospheric correction algorithms include the Second Simulation of the Satellite Signal in the Solar Spectrum (6S), the Moderate Resolution Atmospheric Transmission (MODTRAN), and the Sentinel-2 Atmospheric Correction (Sen2Cor, used by the Copernicus Ground Segment for Sentinel-2 Level-2A processing).

The calibration makes the imagery's radiometric values comparable across acquisitions and across sensors. A vegetation index computed from uncalibrated digital numbers cannot be compared across acquisitions, because the digital numbers vary with the sensor's calibration, the solar irradiance (which varies with the sun angle, the season, and the latitude), and the atmospheric conditions (which vary with the weather and the aerosol load). The food security domain's crop condition assessments depend on the temporal consistency of the vegetation indices: the specialist compares the current NDVI value with the historical NDVI values for the same parcel and the same growth stage, and the comparison is only meaningful if the NDVI values are computed from calibrated surface reflectance.

A miscalibrated image produces NDVI values that are biased relative to the rest of the time series, and the bias is mistaken for a change in crop condition.

Mosaicking combines multiple overlapping images into a single seamless image, transforming a set of scene-based acquisitions into a continuous coverage of the area of interest. The mosaicking process has four steps. The overlap identification step determines the overlapping areas between adjacent images, using the georeferenced image footprints. The color balancing step adjusts the radiometric values of the overlapping images so that the seams are not visible: the adjustment is based on the histogram matching of the overlapping areas, or on a global color balancing that adjusts all images to a common reference. The seam feathering step creates a gradual transition across the seam, rather than a hard boundary, by blending the pixel values in the overlap area using a distance-weighted average. The mosaic production step produces the final mosaic, with the feathered seams and the color-balanced images.

The mosaicking is particularly important for the flood risk domain, where the inundation maps must cover the entire affected area, and for the energy infrastructure domain, where the change detection must cover the entire monitored corridor. A mosaic that is not color-balanced will show visible seams at the scene boundaries, and the seams will be mistaken for real features by the change detection algorithms, producing false positives. A mosaic that is not seam-feathered will show hard boundaries at the scene edges, and the boundaries will appear as linear features in the analysis.

The United States Geological Survey's Landsat Analysis Ready Data (ARD) product provides calibrated surface reflectance at the scene level, tiled to a standard grid (the Worldwide Reference System-2), allowing consumers to mosaic the scenes themselves. The European Space Agency's Sentinel-2 Level-2A product provides calibrated surface reflectance at the scene level, with the Sen2Cor atmospheric correction applied by the Copernicus Ground Segment. The geospatial office uses these products through the processing engine, with the data pipeline operator configuring the mosaicking parameters (the color balancing method, the seam feathering width, the output extent) and the technical lead validating the mosaic against the source scenes (checking for visible seams, color inconsistencies, and geometric misalignment at the scene boundaries).

9.4 Projection Management

Projection management maintains the coordinate reference systems and map projections used in the office's work, and transforms data between coordinate systems and projections as required by the analysis. The coordinate reference system is the mathematical framework that relates coordinates to positions on the Earth's surface. The map projection is the transformation that converts the three-dimensional Earth's surface to a two-dimensional representation. Every map projection introduces distortions in area, shape, distance, and direction, and the choice of projection determines which distortions are minimized and which are accepted.

The office maintains a standard coordinate reference system for its work. The standard CRS includes a datum (the mathematical model of the Earth's shape: WGS 84, GRS 80, or a local datum), a geographic coordinate system (latitude and longitude, expressed in degrees), and a projected coordinate system (a planar coordinate system in metres, derived from the geographic coordinate system through a map projection). The standard projected CRS is typically the Universal Transverse Mercator (UTM) zone for the office's region, or a national projection such as the British National

Grid (EPSG:27700, Transverse Mercator on the Airy 1830 ellipsoid), the Lambert-93 for France (EPSG:2154, Lambert Conformal Conic on the GRS 80 ellipsoid), or the Singapore Transverse Mercator (EPSG:3414, Transverse Mercator on the WGS 84 ellipsoid).

The UTM system divides the Earth into 60 zones, each 6 degrees of longitude wide, with a central meridian at the zone's centre. The UTM projection is a Transverse Mercator that preserves shape and distance within each zone but introduces distortion at the zone boundaries. The distortion is less than 1 part in 1,000 within the zone, which is sufficient for most applications. An office that analyzes flooding across a region that straddles two UTM zones must either reproject the data to one zone and accept the distortion at the edge of the region, or use a different projection that covers the entire region (such as an Albers Equal Area Conic, which preserves area across a wide east-west extent but distorts shape).

The choice of projection has consequences for the accuracy of the analysis. Areal calculations performed on a shape-preserving projection (such as UTM or Lambert Conformal Conic) will return slightly inaccurate areas, because the projection does not preserve area. The error is small within a UTM zone (less than 0.1 per cent) but increases at the zone boundaries. For applications where the area accuracy is critical (such as the crop yield forecast, where the area of each crop type is multiplied by the per-hectare yield), the office should use an equal-area projection (such as Albers Equal Area or Mollweide) or apply an area correction factor. Distance calculations performed on an equal-area projection will return inaccurate distances, because the projection does not preserve distance. For applications where the distance accuracy is critical (such as the buffer analysis around an infrastructure corridor), the office should use an equidistant projection or perform the distance calculation on the geographic coordinates using the great-circle distance formula.

The United States Geological Survey maintains the National Spatial Reference System (NSRS) as the standard coordinate reference system for the Survey's work. The NSRS is maintained by the National Geodetic Survey (NGS), which provides the coordinate transformation tools through the NGS Coordinate Conversion and Transformation Tool (NCAT). The European Space Agency's Copernicus programme maintains its data in WGS 84 and provides the transformation tools through the Sentinel Application Platform (SNAP). The national mapping agencies maintain their national coordinate reference systems and provide transformation tools through their spatial data infrastructures.

9.5 Cloud Masking

Cloud masking identifies clouds and cloud shadows in optical imagery and masks them out of the analysis. The process is critical for the food security and energy infrastructure domains, which depend on optical imagery and which operate in regions and seasons where cloud cover is common. A vegetation index computed over a cloud-contaminated pixel will produce a falsely low value (clouds are bright in the visible bands, which the NDVI formula interprets as low vegetation vigour), and a change detection algorithm will identify the cloud or cloud shadow as a change, producing a false positive.

Clouds are detected through spectral tests: clouds are typically bright in the visible bands (red, green, blue) and dark in the thermal bands (if the sensor has thermal bands, as Landsat does through its Thermal Infrared Sensor). Cloud shadows are detected through geometric tests: a cloud shadow is

located at a predictable offset from the corresponding cloud, based on the solar geometry (the sun elevation and azimuth angles at the time of acquisition, which are recorded in the imagery's metadata). The shadow's offset is calculated as: the cloud's height (estimated from the cloud's brightness temperature in the thermal band) multiplied by the tangent of the solar zenith angle, in the direction of the solar azimuth. The algorithms combine these spectral and geometric tests with temporal tests that compare the imagery with a cloud-free reference (a composite of cloud-free pixels from previous acquisitions).

The Landsat Collection 2 Level-2 product provides cloud masking through the Quality Assessment (QA) band, which identifies clouds, cloud shadows, snow, water, and other quality issues at the pixel level. The QA band uses a bit-packed encoding, where each bit of a 16-bit integer represents a specific quality flag. The consumer unpacks the QA band to extract the individual flags. The Sentinel-2 Level-2A product provides cloud masking through the Scene Classification Layer (SCL), which identifies clouds (opaque and cirrus), cloud shadows, snow, water, vegetation, bare soil, and other classes at 20-metre resolution. The SCL is produced by the Sen2Cor processor as part of the Level-2A atmospheric correction.

The accuracy of the cloud mask determines the accuracy of the analysis. A cloud mask that under-detects clouds (false negatives) will leave cloud-contaminated pixels in the analysis, producing falsely low vegetation index values and false change detections. A cloud mask that over-detects clouds (false positives) will mask legitimate pixels, reducing the data available for the analysis and producing data gaps. The balance between false negatives and false positives is adjusted through the confidence threshold: a lower threshold detects more clouds (fewer false negatives, more false positives), while a higher threshold detects fewer clouds (more false negatives, fewer false positives). The office's standard threshold is set by the technical lead, based on the analysis's tolerance for cloud contamination: the food security domain's crop condition assessment is more tolerant of cloud contamination (because the assessment is based on a time series, and a single cloud-contaminated pixel can be identified and excluded), while the energy infrastructure domain's change detection is less tolerant (because a single false change detection triggers an alert that the consumer must investigate).

The challenge of cloud masking is acute in tropical and subtropical regions, where persistent cloud cover makes it difficult to acquire cloud-free optical imagery. In these regions, the office may use synthetic aperture radar (SAR) imagery, which penetrates clouds, as a complement to the optical imagery. The Sentinel-1 C-band SAR provides imagery regardless of cloud cover, and the imagery can be used for flood extent mapping (standing water appears dark in SAR imagery because the smooth water surface reflects the radar pulse away from the sensor), for crop monitoring (the radar backscatter is sensitive to the crop's structure and moisture), and for change detection (the radar detects changes in the surface roughness). The data fusion of the optical and radar imagery is the mechanism by which the office maintains its analytical capability through the cloud cover, and the cloud masking is the preprocessing step that makes the fusion effective.

9.6 Data Fusion

Data fusion combines data from multiple sources into a single integrated product that is more informative than any of the individual sources. The fusion is conducted at three levels. Pixel-level

fusion combines the data at the resolution of the finest source: a pixel-level fusion of optical imagery and radar imagery produces a fused image that has the spectral richness of the optical imagery and the cloud-penetration of the radar imagery. Feature-level fusion combines the features extracted from different sources into a unified feature set: a feature-level fusion of optical features (vegetation indices, texture measures) and radar features (backscatter coefficients, coherence) produces a unified feature set that can be classified more accurately than either set alone. Decision-level fusion combines the decisions made from different sources: a decision-level fusion of an optical-based crop classification and a radar-based crop classification combines the strengths of the two classifications.

The four domains benefit from the fusion of different data sources. The flood risk domain benefits from the fusion of optical imagery (for land cover and flood extent in cloud-free conditions), radar imagery (for flood extent under clouds), and elevation data (for the topography that the hydraulic model uses). The food security domain benefits from the fusion of optical imagery (for the vegetation indices), weather data (for the growing conditions), and soil moisture data (for the root zone water availability, which can be estimated from the SMAP (Soil Moisture Active Passive) satellite or from the Sentinel-1 SAR backscatter). The energy infrastructure domain benefits from the fusion of high-resolution optical imagery (for the change detection), elevation data (for the line-of-sight analysis and the slope stability assessment), and the infrastructure asset registry (for the asset locations). The minerals domain benefits from the fusion of geological data (for the formation), geophysical data (for the subsurface structure: aeromagnetic for magnetic anomalies, gravity for density variations), and geochemical data (for the element concentrations).

The fusion algorithms vary by level. Pixel-level fusion algorithms include the Intensity-Hue-Saturation (IHS) transform (which replaces the intensity component of a low-resolution multispectral image with a high-resolution panchromatic image, producing a fused image with the spectral content of the multispectral image and the spatial detail of the panchromatic image), the principal component analysis (PCA) fusion (which replaces the first principal component of a low-resolution multispectral image with a high-resolution panchromatic image), and the wavelet fusion (which decomposes the images into wavelet components and combines the high-frequency components from the high-resolution image with the low-frequency components from the low-resolution image). Feature-level fusion algorithms include the concatenation of feature vectors (which simply combines the features from different sources into a single vector) and the feature selection (which selects the most informative features from the combined set). Decision-level fusion algorithms include the majority voting (which takes the majority decision from multiple classifiers), the weighted average (which takes the weighted average of the classifiers' confidence scores), and the Dempster-Shafer evidence theory (which combines the classifiers' decisions using a formal framework for reasoning under uncertainty).

The United States Geological Survey integrates Landsat optical imagery with elevation data from the 3D Elevation Program, radar data from the Sentinel-1 programme, and other sources through the Survey's data integration tools, which are provided through the EarthExplorer platform and the USGS data portals. The European Space Agency's Copernicus programme provides data fusion through the Copernicus Data and Information Access Services (DIAS), which integrate Sentinel data with other sources through cloud-based processing platforms (the five DIAS platforms: CREODIAS, MUNDI, ONDA, Sobloo, and WEkEO). The geospatial office maintains the capability

to fuse data from multiple sources through the processing engine, with the data pipeline operator configuring the fusion parameters and the technical lead validating the fused product against the source data.

9.7 The Geometric Integrity Guarantee

The geometric integrity of a processed product is established through validation against independent check points (ICPs) — points that are not used in the georeferencing transformation, and that therefore test the transformation's accuracy across the image rather than at the control points. The validation is conducted by the data pipeline operator on every processed product, and the technical lead reviews the validation results on a sample basis. The validation compares the processed product's geometry with the independent check points, calculates the root mean square error (RMSE), and verifies that the error is within the specified tolerance.

The RMSE is calculated as: RMSE = sqrt((1/n) × Σ(x_predicted - x_observed)² + (y_predicted - y_observed)²), where n is the number of check points, x_predicted and y_predicted are the coordinates predicted by the transformation, and x_observed and y_observed are the observed coordinates. The RMSE is typically reported at the 95 per cent confidence level (which is 1.96 times the standard deviation for a normally distributed error), and the tolerance is specified in the processing standards. The product that passes the validation (RMSE ≤ tolerance) is released to the modeling workbench; the product that fails (RMSE > tolerance) is reprocessed with corrected parameters or rejected.

The geometric accuracy of the processed products determines the geometric accuracy of the models and the products, and the geometric accuracy of the products determines the credibility of the office's intelligence. A flood inundation map that is geometrically inaccurate will point at the wrong parcels, and the evacuation guidance will be sent to the wrong households. A change detection map that is geometrically inaccurate will identify changes where there are none, and the maintenance advisories will be issued for assets that have not changed. A mineral prospectivity map that is geometrically inaccurate will identify targets at the wrong coordinates, and the exploration licenses will be applied for in the wrong locations.

The validation is conducted for every processed product, without exception. The pressure to waive the validation will arise, particularly during flood events when the demand for intelligence is urgent, but the validation is the operational practice that prevents the office from releasing products that point at the wrong place on the ground. The Landsat Collection 2 Level-1 product requires an RMSE of less than 12 metres. The Sentinel-2 Level-1C product requires a sub-pixel accuracy (less than 10.3 metres at 95 per cent confidence). The geospatial office's processing standards should specify the RMSE tolerance for each product type, and the tolerance should be at least as stringent as the source data's specified accuracy.

9.8 Worked Example: Processing Sentinel-2 Imagery for a Crop Condition Assessment

The data pipeline operator receives a Sentinel-2 Level-1C image acquired over the agricultural region at 06:00 UTC, with 15 per cent cloud cover. The image was acquired by Sentinel-2A (the acquisition identifier is

S2A_MSIL1C_20240116T060000_N0509_R034_T42SXA_20240116T080000). The operator ingests the image into the processing engine under the public-service charge code (PUB-2024-120) and records the provenance: data source (Sentinel-2A, acquisition ID as above), acquisition method (subscription, Copernicus free and open data), acquisition date (2024-01-16 06:00 UTC), ingestion date (2024-01-16 08:15 UTC), license terms (Copernicus free and open, attribution required), intended domain consumer (food security).

The operator configures the processing run. The georeferencing is not required, because the Sentinel-2 Level-1C product is already georeferenced by the Copernicus Ground Segment to sub-pixel accuracy. The calibration is configured for surface reflectance processing: the operator selects the Sen2Cor processor (version 2.11) to perform the atmospheric correction, converting the Level-1C TOA reflectance to Level-2A surface reflectance. The Sen2Cor processor requires the auxiliary atmospheric data (aerosol optical thickness, water vapor, ozone), which it estimates from the imagery itself through the Dark Object Subtraction method and the cirrus band (band 10, 1.37 µm, which is sensitive to cirrus clouds).

The cloud masking is configured to use the Scene Classification Layer (SCL) produced by the Sen2Cor processor, at confidence threshold 60 per cent. The SCL identifies 15 per cent cloud cover (matching the acquisition log), distributed in the northwestern corner of the image. The operator visually inspects the SCL against the true-colour image and confirms that the cloud mask is accurate: the masked areas correspond to visible clouds, and no legitimate pixels are masked.

The mosaicking is not required, because the image is a single scene covering the entire agricultural region. The projection management is configured to transform the data from the native UTM Zone 42N (EPSG:32642) to the office's standard projection, which is also UTM Zone 42N (no transformation required).

The operator submits the processing run to the processing engine. The Sen2Cor processor runs for approximately 12 minutes on the office's processing server, producing the Level-2A surface reflectance product with the SCL cloud mask. The operator reviews the processing log: the atmospheric correction completed successfully, the surface reflectance values are within the expected range (0 to 0.4 for visible bands, 0 to 0.6 for near-infrared band), and no errors or warnings were reported.

The technical lead validates the processed product. The geometric validation is not required (the product is already georeferenced by the Copernicus Ground Segment), but the technical lead verifies the geometric accuracy by checking 5 independent check points against the cadastre: all 5 check points are within 8 metres of the expected position, which is within the Sentinel-2 specification (10.3 metres at 95 per cent confidence). The radiometric validation is conducted by comparing the surface reflectance values at a stable reference target (a bare soil area that has been used as a validation site for the previous 3 years): the surface reflectance values are within 2 per cent of the historical mean, which is within the expected variability. The cloud mask validation is conducted by visually inspecting 20 randomly selected 1-kilometre tiles: the cloud mask accurately identifies clouds in 19 of 20 tiles, with 1 false positive (a bright agricultural field that the SCL misclassified as cloud). The technical lead records the validation results in the provenance record.

The technical lead releases the processed product to the modeling workbench. The resource economics specialist retrieves the surface reflectance bands (blue, green, red, near-infrared, red-edge,

shortwave infrared), the SCL cloud mask, and the provenance record from the workbench. The specialist configures the crop model plug-in to calculate the NDVI (from the red and near-infrared bands: NDVI = (NIR - Red) / (NIR + Red)), the NDWI (from the near-infrared and shortwave infrared bands: NDWI = (NIR - SWIR) / (NIR + SWIR)), and the EVI (from the blue, red, and near-infrared bands: EVI = G × (NIR - Red) / (NIR + C1 × Red - C2 × Blue + L), where G = 2.5, C1 = 6, C2 = 7.5, L = 1). The specialist stratifies the indices by parcel using the cadastre feed, producing the parcel-level condition assessments for the weekly crop condition bulletin.

This worked example traces the processing capability's six sub-capabilities through a single processing run. The georeferencing was provided by the data provider (Copernicus Ground Segment). The calibration was conducted by the Sen2Cor processor (TOA to surface reflectance). The mosaicking was not required (single scene). The projection management confirmed the native projection matched the office's standard (no transformation). The cloud masking was conducted by the SCL at 60 per cent confidence threshold. The data fusion was not required for this run (single-source optical imagery). The geometric integrity was validated against 5 independent check points. The provenance record was maintained throughout, from the acquisition ID through the processing parameters to the validation results, and the chain is unbroken. The processed product is now ready for the modeling workbench, where the resource economics specialist will use the surface reflectance data to produce the weekly crop condition bulletin for the strategic reserve authority.

Chapter 10 — Modeling

10.1 The Modeling Function

Modeling is the fourth capability of the pipeline. It is the domain-specific analytical layer — the only pipeline capability that varies by domain, varying through the mechanism of swappable modeling heads, treated in Chapter 6. The modeling function is conducted through the modeling workbench submodule of the geospatial information system, which hosts the domain plug-ins that implement the domain-specific methods. The workbench provides the shared environment, the data interface, the parameter management, the validation framework, and the provenance recording; the plug-ins provide the analytical methods.

The modeling function is operated by the specialists, with the technical lead configuring the workbench and validating the model runs, and the office manager overseeing the production. The hazard specialist operates the hydrology and change detection plug-ins; the resource economics specialist operates the crop model and prospectivity plug-ins. The specialists are the office's domain experts, and their expertise is the analytical foundation of the office's intelligence. The specialists' work is specified in the job breakdown sheets in Chapters 17 and 18.

The modeling workbench enforces a separation between the shared environment and the plug-ins. The shared environment handles the data ingestion, the parameter management, the validation framework, and the provenance recording; the plug-ins handle the domain-specific methods. A new domain can be added by developing a new plug-in, without modifying the shared environment; an existing plug-in can be updated without affecting the other plug-ins. The separation allows the office to add domains without rebuilding the workbench, and allows the plug-in developers to iterate on the domain-specific methods without affecting the rest of the pipeline.

The four plug-ins implement the domain-specific analytical methods that produce the office's intelligence. Each plug-in consumes processed data from the processing engine and base layers from the spatial data infrastructure, applies the domain-specific model, and produces the analytical output that the production function transforms into the final deliverable. The plug-ins are treated in sections 10.2 through 10.5. Model validation (section 10.6) and model provenance (section 10.7) are the cross-cutting concerns that apply to all four plug-ins.

10.2 The Hydrology Plug-In

The hydrology plug-in implements the flood risk models. The hazard specialist operates it for the flood risk domain. The plug-in comprises three coupled models: the hydrological model that simulates the rainfall-runoff process, the hydraulic model that simulates the flow of water in channels and floodplains, and the inundation model that maps the extent and depth of flooding.

The hydrological model simulates the conversion of rainfall into runoff, accounting for infiltration (the water that soaks into the soil), evapotranspiration (the water that evaporates from the surface and transpires through vegetation), and subsurface flow (the water that moves through the soil and eventually reaches the channel as baseflow). The hydrological model takes the rainfall data from the rain gauge network or the weather radar, the soil data from the soil map, the land cover data from the land cover base layer, and the topographic data from the digital elevation model, and produces a

hydrograph — a time series of the river flow at the catchment outlet or at specified points along the channel network.

The rainfall-runoff transformation is the core of the hydrological model. Several approaches exist. The Sacramento Soil Moisture Accounting (SAC-SMA) model, used by the United States National Weather Service for the National Water Model, divides the soil into two zones (an upper zone representing the topsoil and a lower zone representing the subsoil) and tracks the moisture in each zone through a set of parameters that represent the infiltration capacity, the percolation rate, and the baseflow recession. The Green-Ampt infiltration model, used in many distributed hydrological models, computes the infiltration rate from the soil's hydraulic conductivity, the capillary pressure head, and the soil's initial moisture content. The Soil Conservation Service (SCS) Curve Number method, used in many engineering applications, estimates the runoff from the rainfall and a curve number that represents the catchment's runoff characteristics (which depend on the soil type, the land cover, and the antecedent moisture condition).

The hydraulic model simulates the flow of the runoff through the channel network and the floodplain. The model takes the hydrograph from the hydrological model as its upstream boundary condition, the river stage from the sensor feeds as its downstream boundary condition, and the channel geometry from the hydrography base layer and the DEM. The model solves the one-dimensional or two-dimensional Saint-Venant equations (the conservation of mass and the conservation of momentum for open-channel flow) to compute the water surface elevation, the flow velocity, and the flow depth at each point in the channel and the floodplain. The one-dimensional model (such as HEC-RAS 1D, developed by the US Army Corps of Engineers Hydrologic Engineering Center) represents the channel as a series of cross-sections and solves the equations along the channel's longitudinal axis; the model is computationally efficient but cannot represent the lateral flow in the floodplain. The two-dimensional model (such as HEC-RAS 2D, MIKE 21, or TUFLOW) represents the channel and the floodplain as a two-dimensional grid and solves the equations in both the longitudinal and the lateral directions; the model is computationally intensive but can represent the complex flow patterns in the floodplain, including the flow around buildings, levees, and road embankments.

The inundation model maps the extent and depth of the flooding by intersecting the hydraulic model's water surface elevation with the digital elevation model. The intersection produces a polygon representing the flooded area and a raster representing the flood depth (the water surface elevation minus the ground elevation). The flood depth raster is the basis for the exposure assessment: the raster is overlaid with the cadastre to identify the affected parcels, with the population data to estimate the affected population, and with the infrastructure asset registry to identify the affected infrastructure.

The United States National Weather Service operates the National Water Model (NWM), a continental-scale hydrological modelling framework that became operational in 2016. The NWM uses the WRF-Hydro modelling system (Weather Research and Forecasting model with Hydrology), which couples the Noah-MP land surface model with a distributed routing model. The NWM produces streamflow forecasts at approximately 2.7 million river locations across the United States, with a 1-hour analysis and assimilation cycle, a 10-day deterministic forecast, and a 30-day ensemble forecast. The NWM is calibrated against the United States Geological Survey's stream gauge network (approximately 8,500 gauges) and is used by the National Weather Service's River Forecast Centers

for flood forecasting. The European Flood Awareness System (EFAS), operational since 2005 and managed by the European Commission's Joint Research Centre, uses the LISFLOOD hydrological model to produce pan-European flood forecasts with a 3-to-10-day lead time, calibrated against the national hydrological services' gauge networks.

The geospatial office maintains its hydrological models (typically a distributed model such as HEC-HMS for the rainfall-runoff and HEC-RAS 2D for the hydraulics, or a coupled model such as MIKE SHE/MIKE 11), calibrates them against observed flood events recorded by the national hydrological service's gauge network, and uses them to produce flood inundation maps and exposure reports. The calibration adjusts the model parameters (the Manning's n roughness coefficients, the infiltration parameters, the baseflow recession constants) until the model's simulated hydrograph matches the observed hydrograph at the gauge stations. The validation compares the model's inundation extent with the observed extent (from satellite imagery or field surveys) and the model's inundation depth with the observed depth (from high-water marks or gauge records). The validation is the operational practice that establishes the analytical integrity of the flood risk domain.

10.3 The Crop Model Plug-In

The crop model plug-in implements the food security models. The resource economics specialist operates it for the food security domain. The plug-in comprises three components: the crop condition model that assesses the state of the crops from satellite imagery, the yield model that estimates the crop yield from the condition and the historical data, and the procurement advisory model that synthesizes the yield forecast with the reserve stock levels to produce a release or buy signal.

The crop condition model computes vegetation indices from the calibrated surface reflectance. The Normalized Difference Vegetation Index (NDVI), proposed by Rouse et al. in 1973, is the ratio of the difference between the near-infrared and the red reflectance to their sum: NDVI = (NIR – Red) / (NIR + Red). The NDVI exploits the fact that healthy vegetation absorbs red light (for photosynthesis) and reflects near-infrared light (due to the internal structure of the leaf): a healthy crop has a high NDVI (0.6 to 0.8), while bare soil has a low NDVI (0.1 to 0.2), and water has a negative NDVI. The Enhanced Vegetation Index (EVI), proposed by Huete et al. in 2002, was developed to correct for the atmospheric scattering and the canopy background effects that limit the NDVI in dense vegetation: EVI = G × (NIR – Red) / (NIR + C1 × Red – C2 × Blue + L), where G = 2.5, C1 = 6, C2 = 7.5, and L = 1. The EVI is more sensitive to high-biomass vegetation than the NDVI and does not saturate as quickly. The Leaf Area Index (LAI) is the ratio of the one-sided leaf area to the ground area, and it is estimated from the vegetation indices through empirical relationships that are calibrated for each crop type.

The crop condition model compares the current vegetation index value with the historical distribution for the same parcel and the same phenological stage. The phenological stage (emergence, tillering, stem elongation, heading, flowering, grain filling, maturity) is determined from the crop calendar (which specifies the expected planting and harvest dates for each crop and each region) and from the accumulated growing degree days (the cumulative temperature above a base temperature, which drives the crop's development). The comparison produces a condition score: the current NDVI is compared with the historical NDVI distribution for the same parcel, the same crop, and the same phenological stage, and the condition score is expressed as a percentile (e.g., "the current NDVI is at

the 75th percentile of the historical distribution, indicating above-normal condition") or as a departure from the median (e.g., "the current NDVI is 15 per cent above the historical median").

The yield model combines the condition score with the weather data, the soil moisture data, and the historical yield data to produce a yield forecast. The yield model may be statistical (a regression model that relates the NDVI at a specific phenological stage to the final yield, calibrated using historical data) or mechanistic (a crop growth model that simulates the crop's physiological processes, such as the World Food Studies (WOFOST) model developed by Wageningen University, or the Decision Support System for Agrotechnology Transfer (DSSAT) developed by the University of Georgia). The statistical model is simpler and requires less data but is less transferable (the regression coefficients are specific to the crop, the region, and the calibration period). The mechanistic model is more complex and requires more data (weather, soil, crop management) but is more transferable and can simulate the effects of management changes (irrigation, fertilization) on the yield.

The procurement advisory model combines the yield forecast with the strategic reserve authority's stock levels and consumption rates to produce a release or buy signal. The yield forecast is converted to a production estimate (the yield multiplied by the planted area, which is obtained from the cadastre or from the agricultural ministry's planted area survey), and the production estimate is compared with the expected consumption (the population multiplied by the per-capita consumption rate, adjusted for trends). If the production estimate exceeds the consumption (a surplus), the advisory recommends a release from the reserve; if the production estimate is below the consumption (a deficit), the advisory recommends a buy for the reserve.

The GEOGLAM initiative — the Group on Earth Observations Global Agricultural Monitoring, launched in 2011 by the G20 Agricultural Ministers — coordinates the global crop monitoring effort and provides the framework for the national crop monitoring services. GEOGLAM's Crop Monitor provides a monthly global crop condition assessment that synthesizes the assessments of 35 participating organizations. The European Commission's Mars (Monitoring Agricultural Resources) bulletin, produced by the Mars Unit of the Joint Research Centre since 1992, provides crop condition assessments and yield forecasts for the European Union and other regions, using the WOFOST crop growth model and the MARS (Monitoring Agriculture with Remote Sensing) database of weather stations and crop observations. The United States Department of Agriculture's Foreign Agricultural Service (FAS) provides crop condition assessments and yield forecasts through its World Agricultural Production (WAP) circular, using a combination of satellite imagery analysis, crop models, and agricultural attaché reports from US embassies.

The geospatial office maintains its crop models (typically a statistical NDVI-yield regression model for the weekly bulletin and a WOFOST or DSSAT mechanistic model for the seasonal forecast), calibrates them against harvest outcomes reported by the agricultural ministry's statistical service, and uses them to produce crop condition assessments, yield forecasts, and procurement advisories. The calibration adjusts the regression coefficients (for the statistical model) or the crop parameters (for the mechanistic model) until the model's yield forecast matches the reported harvest. The validation compares the model's yield forecast with the reported harvest for the current season, and the comparison is recorded in the provenance chain.

10.4 The Change Detection Plug-In

The change detection plug-in implements the energy infrastructure models. The hazard specialist operates it for the energy infrastructure domain. The plug-in comprises the change detection algorithms that identify changes in the imagery between acquisitions, the encroachment detection algorithms that identify third-party construction or vegetation growth near the monitored assets, and the vulnerability scoring algorithms that assess the assets' vulnerability to the detected changes.

The change detection algorithm compares two acquisitions of the same area and identifies the pixels or objects that have changed. Two approaches exist. Image-to-image change detection compares the radiometric values of the two acquisitions: the image differencing subtracts the radiometric values of the earlier image from those of the later image and identifies the pixels where the difference exceeds a threshold; the image ratioing divides the later image by the earlier image and identifies the pixels where the ratio deviates from 1; the change vector analysis computes the direction and magnitude of the change in the multispectral feature space and identifies the pixels where the magnitude exceeds a threshold. Object-to-object change detection first segments the imagery into objects (contiguous regions of similar spectral characteristics) and then compares the objects between acquisitions: an object that has changed in area, shape, or spectral characteristics is identified as a change. The object-based approach is more robust to geometric misregistration and to radiometric differences between acquisitions, because the comparison is at the object level rather than the pixel level, but it requires the segmentation step, which adds complexity and processing time.

The encroachment detection algorithm applies the change detection to a buffer around each asset in the infrastructure asset registry. The buffer width is determined by the asset's type: a transmission line may have a buffer of 30 metres (the vegetation clear zone), a substation may have a buffer of 100 metres (the security perimeter), a pipeline may have a buffer of 15 metres (the right-of-way). The algorithm identifies the changes within the buffer and classifies them by type: vegetation encroachment (the vegetation has grown into the buffer), third-party construction (a new structure has been built within the buffer), ground subsidence (the ground elevation has decreased, indicating potential slope instability), and excavation (the ground surface has been disturbed, indicating potential damage to buried infrastructure).

The vulnerability scoring algorithm combines the detected changes with the asset's criticality and the change's proximity to assess the asset's vulnerability. The asset's criticality is determined by the asset registry: a transmission line serving 50,000 customers is more critical than one serving 5,000; a substation that is a grid interconnection point is more critical than one that is a local distribution point. The change's proximity is the distance from the change to the asset: a change within 5 metres of a transmission tower is more urgent than a change at 25 metres. The vulnerability score is the product of the criticality and the proximity (adjusted for the change type), and the score prioritizes the maintenance advisories: the assets with the highest scores are the ones that the energy sector entity should address first.

The European Maritime Safety Agency (EMSA) operates the CleanSeaNet service, which uses satellite imagery to detect oil spills and monitor maritime infrastructure. The service processes Sentinel-1 radar imagery and high-resolution optical imagery and provides the change detection results to the member states' coast guards. The United States Department of Homeland Security's

Critical Infrastructure Protection programme, established under the Homeland Security Presidential Directive 7 (HSPD-7) in 2003, uses commercial satellite imagery to monitor pipelines, refineries, and transmission lines through change detection algorithms. The programme coordinates with the National Geospatial-Intelligence Agency for imagery acquisition and with the sector-specific agencies (the Department of Energy for energy infrastructure, the Transportation Security Administration for pipelines) for the interpretation of the change detection results.

The geospatial office maintains its change detection algorithms, validates them against ground truth inspections conducted by the energy sector entity's field teams, and uses them to produce infrastructure risk maps and encroachment alerts. The validation compares the algorithm's change detections with the field teams' inspection reports: the true positives (changes detected by the algorithm and confirmed by the field team), the false positives (changes detected by the algorithm but not confirmed by the field team), and the false negatives (changes detected by the field team but missed by the algorithm) are counted, and the detection rate (true positives divided by all actual changes) and the false alarm rate (false positives divided by all detected changes) are calculated. The accuracy metrics are recorded in the provenance chain, and the algorithm's parameters (the change threshold, the buffer width, the vulnerability scoring weights) are adjusted to optimize the accuracy.

10.5 The Prospectivity Plug-In

The prospectivity plug-in implements the minerals models. The resource economics specialist operates it for the minerals domain. The plug-in comprises the prospectivity models that assess the mineral potential of a geological formation from the geological, geophysical, and geochemical data, and the resource estimation models that estimate the quantity and quality of the mineral resource.

The prospectivity model combines the geological, geophysical, and geochemical data into a mineral potential map. Two approaches exist. The knowledge-driven approach encodes the geological expert's criteria into a set of weights: the expert identifies the geological features that are associated with the mineral deposit type (e.g., for porphyry copper deposits: intrusive rocks of intermediate composition, hydrothermal alteration zones, geochemical anomalies in copper and molybdenum, magnetic anomalies associated with the intrusive), assigns weights to each feature, and the model combines the weighted features into a prospectivity map. The weights-of-evidence method, developed by Bonham-Carter et al. in 1989, is a knowledge-driven approach that uses Bayes' theorem to compute the weights from the spatial association between the known deposits and the geological features: the positive weight (W+) is the log odds of the deposit occurring given the feature's presence, and the negative weight (W-) is the log odds of the deposit occurring given the feature's absence. The fuzzy logic method, developed by An et al. in 1991, uses fuzzy membership functions to represent the uncertainty in the geological expert's criteria and combines the fuzzy memberships using fuzzy operators (fuzzy AND, fuzzy OR, fuzzy gamma).

The data-driven approach uses known deposits to train a machine learning model. The model is trained on a set of positive examples (the locations of known deposits) and negative examples (the locations where deposits are known not to occur, or randomly selected locations if true negatives are not available), and the model learns the spatial patterns that distinguish the positive from the negative examples. The random forest classifier, developed by Breiman in 2001, is an ensemble of decision trees that is robust to overfitting and can handle high-dimensional data with mixed variable types. The

support vector machine (SVM), developed by Cortes and Vapnik in 1995, finds the hyperplane that maximally separates the positive from the negative examples in the feature space and can handle non-linear separations through the kernel trick. The artificial neural network (ANN), which has been applied to prospectivity mapping since the 1990s, learns the non-linear relationships between the geological features and the deposit occurrence through a network of interconnected nodes.

The resource estimation model combines the prospectivity map with the historical exploration data — drill holes, trench samples, bulk samples — to estimate the quantity and quality of the mineral resource. The estimation is conducted using geostatistical methods: ordinary kriging, which estimates the grade at unsampled locations from the grades at sampled locations, weighted by the spatial covariance of the grades; indicator kriging, which estimates the probability that the grade exceeds a cutoff; and conditional simulation, which generates multiple equally probable realizations of the grade distribution to assess the uncertainty in the estimate. The resource estimation follows international reporting standards: the Joint Ore Reserves Committee (JORC) Code (Australia), the National Instrument 43-101 (Canada), or the Society for Mining, Metallurgy, and Exploration (SME) Standard (United States), which specify the data quality requirements, the estimation methods, and the reporting categories (inferred, indicated, measured resources; probable and proven reserves).

The critical minerals inventory aggregates the resource estimates for the minerals on the government's critical minerals list. The critical minerals list is determined by the government's assessment of the minerals that are essential to the economy and whose supply is at risk: the United States' List of Critical Minerals (50 minerals, established by the Energy Act of 2020 and updated by the USGS in 2022), the European Union's List of Critical Raw Materials (34 materials, established by the European Commission in 2011 and updated in 2023), and the Australian Critical Minerals List (26 minerals, established by the Department of Industry, Science, and Resources). The inventory provides the strategic stockpile authority with the information needed to plan the stockpile's composition and to assess the supply risk for each mineral.

The United States Geological Survey maintains mineral prospectivity models for the United States through its Mineral Resources Program, and publishes the assessments through the National Minerals Information Center and the USMIN Mineral Deposit Database. Geoscience Australia maintains mineral prospectivity models through its Minerals Systems program and the Exploring for the Future initiative, and publishes the assessments through the Australian Mineral Exploration Reviews. The British Geological Survey maintains mineral prospectivity models through its Mineral Resources programme and publishes the assessments through the BGS Mineral Planning Factsheets. In each case, the survey validates its prospectivity models against known deposits recorded in the national mineral occurrence database, and the validation is the operational practice that establishes the analytical integrity of the minerals domain.

10.6 Model Validation

Model validation verifies that the office's models produce accurate results. The validation is conducted through the comparison of the model outputs with independent ground truth, and the comparison is conducted for every model run, with the results recorded in the provenance chain. The validation selects the ground truth, compares the model output with the ground truth, calculates the accuracy metrics, and verifies that the accuracy is within the specified tolerance. The model run that

passes the validation is released to the production function; the model run that fails is re-run with adjusted parameters or rejected.

The validation protocols vary by domain, because the ground truth and the accuracy metrics vary by domain. The flood risk domain validates its models against observed flood events recorded by the national hydrological service's gauge network. The accuracy metrics include the inundation extent accuracy (the percentage of the observed inundation extent that the model correctly predicts, and the percentage of the model's predicted extent that is not observed — the true positive rate and the false positive rate), the depth accuracy (the root mean square error of the predicted depth at the gauge locations), and the timing accuracy (the difference between the predicted and the observed time of the flood peak). The tolerance for the inundation extent accuracy is typically 80 per cent true positive rate (the model must correctly predict at least 80 per cent of the observed inundation extent); the tolerance for the depth accuracy is typically 0.3 metres RMSE; the tolerance for the timing accuracy is typically 2 hours.

The food security domain validates its models against harvest outcomes reported by the agricultural ministry's statistical service. The accuracy metrics include the yield forecast accuracy (the percentage difference between the forecast yield and the reported harvest yield) and the condition assessment accuracy (the percentage of parcels where the model's condition score matches the field-reported condition). The tolerance for the yield forecast accuracy is typically 10 per cent (the forecast must be within 10 per cent of the reported yield at the district level); the tolerance for the condition assessment accuracy is typically 75 per cent (the model's condition score must match the field-reported condition for at least 75 per cent of the sampled parcels).

The energy infrastructure domain validates its models against ground truth inspections conducted by the energy sector entity's field teams. The accuracy metrics include the detection rate (true positives divided by all actual changes) and the false alarm rate (false positives divided by all detected changes). The tolerance for the detection rate is typically 85 per cent (the model must detect at least 85 per cent of the actual changes); the tolerance for the false alarm rate is typically 20 per cent (no more than 20 per cent of the model's detected changes should be false positives).

The minerals domain validates its models against known deposits recorded in the geological survey's mineral occurrence database. The accuracy metrics include the prospectivity accuracy (the percentage of known deposits that fall within the model's high-prospectivity areas) and the prediction ratio (the ratio of the known deposit density in the high-prospectivity areas to the known deposit density in the overall study area). The tolerance for the prospectivity accuracy is typically 70 per cent (at least 70 per cent of the known deposits must fall within the high-prospectivity areas); the tolerance for the prediction ratio is typically 3.0 (the known deposit density in the high-prospectivity areas must be at least 3 times the overall density).

10.7 Model Provenance

Model provenance is the record of the model version, the model parameters, the model run date, and the validation results. It is the continuation of the provenance chain from the processing function. The model provenance is recorded in the catalog and archive submodule, linked to the processing provenance and the acquisition provenance, forming the unbroken chain from acquisition through archive. The model provenance allows the office to answer the question "how was this model run

conducted?" for any model run it has ever conducted. No model run is released without a complete provenance record.

The model provenance is maintained through the modeling workbench, which records the model version (e.g., HEC-RAS 2D version 6.3.1, WOFOST version 7.1.3, the random forest's scikit-learn version 1.3.0 and the hyperparameters n_estimators=500, max_depth=10), the model parameters (the Manning's n values, the crop growth parameters, the change detection threshold, the prospectivity weights), the run date (the date and time the model was executed), the run duration (the processing time), the input data (the processed products and the base layers used, with their version numbers), and the validation results (the accuracy metrics and the pass/fail determination). The record is linked to the production record, which records the production parameters and the production date, and to the archive record, which records the product identifier and the retention designation.

The United States Geological Survey's data management policy (Survey Manual Chapter 502.1) requires that every model run be recorded with its version, parameters, and validation results, enforced through the Survey's data management system. The Copernicus programme's data quality framework requires that every model run be traceable to its inputs and its parameters, enforced through the programme's quality assurance system. The model provenance, combined with the processing provenance and the acquisition provenance, forms the complete provenance chain that makes the office's products auditable, reproducible, and defensible.

10.8 Worked Example: A Coupled Hydrological-Hydraulic Model Run for a Flood Event

The hazard specialist configures a coupled hydrological-hydraulic model run for the Riverbourne Basin flood event. The event was triggered by 180 millimetres of rainfall over 48 hours, and the river stage at Gauge RB-04 has risen from 2.1 metres to 3.8 metres, crossing the 5-year flood threshold.

The specialist configures the hydrological model (HEC-HMS version 4.10) with the event parameters. The rainfall data is from the 15 rain gauges in the basin, with the data interpolated to a 1-kilometre grid using the inverse distance weighting method. The soil data is from the national soil map (1:50,000 scale), with the soil types classified by their hydrologic soil group (A: high infiltration, B: moderate, C: slow, D: very slow). The land cover data is from the national land cover database (30-metre resolution), with the land cover types classified by their SCS curve number. The curve numbers range from 30 (forest on A-group soil) to 95 (urban on D-group soil). The antecedent moisture condition is set to II (average), based on the rainfall in the preceding 5 days. The loss method is the SCS curve number method, the transform method is the SCS unit hydrograph (with a time of concentration of 4.5 hours for the main channel), and the baseflow method is the exponential recession (with a baseflow recession constant of 0.85 and a ratio-to-peak of 0.15). The hydrological model produces a hydrograph at the catchment outlet: the peak flow is 420 cubic metres per second, occurring 6 hours after the rainfall peak.

The specialist configures the hydraulic model (HEC-RAS 2D version 6.3.1) with the hydrograph as the upstream boundary condition and the observed river stage at Gauge RB-08 (the downstream gauge) as the downstream boundary condition. The terrain data is the LiDAR-derived DEM at 2-metre resolution (the LiDAR was acquired in 2019 and has a vertical RMSE of 0.12 metres). The

mesh is a 5-metre grid (the LiDAR is resampled from 2 metres to 5 metres to reduce the computation time), with breaklines along the levees and the road embankments (to ensure that these features are represented in the mesh). The Manning's n roughness coefficients are assigned by land cover type: 0.035 for the main channel, 0.060 for the urban areas, 0.080 for the forested floodplain, 0.030 for the agricultural floodplain. The hydraulic structures (2 bridges and 1 weir) are represented in the model with their surveyed geometry. The simulation period is 48 hours, with a time step of 10 seconds (to ensure numerical stability). The model run takes 45 minutes on the office's processing server.

The specialist validates the model run against the observed data. The peak stage at Gauge RB-04 (the gauge that triggered the event response) is predicted at 3.72 metres; the observed peak is 3.85 metres. The difference is 0.13 metres, which is within the 0.30-metre tolerance. The peak stage at Gauge RB-06 (a mid-basin gauge) is predicted at 2.91 metres; the observed peak is 2.78 metres. The difference is 0.13 metres, within tolerance. The peak stage at Gauge RB-08 (the downstream gauge) is predicted at 1.94 metres; the observed peak is 1.88 metres. The difference is 0.06 metres, within tolerance. The timing of the peak at all three gauges is within 30 minutes of the observed timing, within the 2-hour tolerance. The specialist compares the model's inundation extent with the Sentinel-1 SAR-derived flood extent (the SAR image was acquired at 05:30 UTC during the event): the true positive rate is 87 per cent (the model correctly predicts 87 per cent of the observed inundation extent), within the 80 per cent tolerance. The false positive rate is 12 per cent (12 per cent of the model's predicted extent is not observed in the SAR), within the 20 per cent tolerance.

The specialist generates the flood inundation map and the exposure report. The inundation map shows the flood extent and depth at 5-metre grid resolution. The exposure report identifies 347 affected parcels (23 residential, 4 commercial, 320 agricultural), estimates 1,200 people within the inundated area (from the census data), and identifies 2 substations and 1 transmission corridor within the inundated area (from the infrastructure asset registry). The specialist documents the model run in the provenance record: HEC-HMS version 4.10 with the SCS curve number method and the exponential recession baseflow; HEC-RAS 2D version 6.3.1 with the 5-metre mesh and the LiDAR DEM; the Manning's n values by land cover; the 48-hour simulation with a 10-second time step; the validation results (peak stage within 0.13 metres at 3 gauges, timing within 30 minutes, inundation extent 87 per cent true positive rate). The provenance record is linked to the processing provenance (the LiDAR DEM and the land cover data) and to the acquisition provenance (the Sentinel-1 SAR image used for the validation). The chain is unbroken from the sensor feeds through the hydrological model, the hydraulic model, the inundation map, and the exposure report, to the archive.

Chapter 11 — Production and Cartography

11.1 The Production Function

Production and cartography is the fifth capability of the pipeline. It transforms model outputs into final deliverables through cartographic design, symbology application, classification marking, and format conversion. The capability is shared across all domain applications, operated by the production operator, and reviewed by the office manager for compliance with cartographic standards and for classification correctness before any product is released. The production function is the last capability before the product reaches the consumer. A product that is analytically correct but cartographically poor is read with difficulty; a product that is analytically correct but misclassified is handled at the wrong security level. Both failures fall on the consumer, and both fall on the office's credibility.

The production function is conducted through the product studio submodule of the geospatial information system, treated in Chapter 29. The product studio manages the cartographic templates, the symbology libraries, the classification marking tools, and the format conversion tools. It provides the interface through which the production operator produces the products, the office manager reviews them, and the specialists provide the analytical content. The design and use of the product studio are specified in the job breakdown sheets in Chapter 16.

The production function operates under two modes. Scheduled production produces the weekly crop condition bulletin, the seasonal yield forecast, the monthly infrastructure risk map, and the annual standing minerals products on a regular cycle. The production operator knows the schedule in advance, prepares the templates, and produces the products on the established deadlines. Event-driven production produces the flood inundation map and the post-event infrastructure damage assessment on demand, in response to a flood event or an infrastructure incident. The two modes share the same production pipeline, but the event-driven mode operates under compressed time constraints, and the production operator must apply the same standards under tighter deadlines. The operator who maintains the standards under event-driven pressure is the operator whose products are trusted by the emergency management authority; the operator who relaxes the standards under pressure produces products that are produced quickly but are not trusted.

The production function has four sub-capabilities, each treated in a separate section of this chapter. Cartographic design (section 11.2) transforms model outputs into maps through the application of cartographic principles. Symbology and classification marking (section 11.3) applies the visual language and the security markings to the product. Format conversion (section 11.4) converts the product from the production format to the formats required by the consumers. The publication pipeline (section 11.5) is the sequence of steps by which a model output is transformed into a released product.

11.2 Cartographic Design

Cartographic design transforms model outputs into maps through the application of cartographic principles. The principles include the selection of the map projection, the design of the symbology, the generalization of the features, the design of the legend, and the layout of the map elements. The

design is the communication design of the office's intelligence. A map that is analytically accurate but cartographically poor will be misread: the consumer will misinterpret the symbology, will overlook the features that are not generalized appropriately for the scale, will not find the information in the legend, and will not be able to locate the area of interest on the map.

The map projection selection determines how the three-dimensional Earth's surface is represented on the two-dimensional map. For the office's products, the projection is typically the office's standard projection (treated in Chapter 9, section 9.4), and the projection is noted on the map. For products that cover a wide area (such as the national flood hazard map), the projection may be an equal-area projection (such as the Albers Equal Area Conic) that preserves the area of the features; for products that cover a narrow area (such as a single-basin flood inundation map), the projection may be the UTM zone that contains the area. The projection note on the map includes the projection name, the datum, and the EPSG code, so that the consumer can reproject the map if needed.

The symbology design assigns symbols, colours, and patterns to the features on the map. The symbology for each product type is specified in the cartographic standards (section 11.6) and applied through the cartographic templates. The flood inundation symbology uses a blue gradient for depth: shallow flooding (0 to 0.5 metres) is light blue, moderate flooding (0.5 to 1.5 metres) is medium blue, deep flooding (over 1.5 metres) is dark blue. The crop condition symbology uses a green-to-red gradient: above-normal condition is dark green, normal condition is light green, below-normal condition is yellow, poor condition is orange, very poor condition is red. The change detection symbology uses red for new features (new construction, new vegetation growth) and blue for removed features (demolished structures, cleared vegetation). The prospectivity symbology uses a warm-colour gradient for high prospectivity (red for the highest, orange for high, yellow for moderate) and a cool-colour gradient for low prospectivity (light blue for low, dark blue for the lowest).

The generalization of the features is the process by which the features are simplified for display at a given scale. Generalization is necessary because the map cannot show every feature at its full detail without becoming cluttered and unreadable. The generalization rules are specified in the standards for each map scale. At 1:25,000, the rules preserve building footprints, road centrelines, and individual watercourses. At 1:100,000, the rules aggregate buildings into built-up areas, simplify the road network to principal routes, and aggregate watercourses into the principal channels. At 1:500,000, the rules reduce the road network to highways and the watercourses to major rivers. The generalization is conducted through the product studio's generalization tools, which apply the rules automatically based on the map scale. The production operator verifies that the generalization is correct by visually inspecting the map at the target scale.

The legend design places the legend in the standard location (typically the lower right corner of the map), lists the symbology in the order of the features' importance, and includes the scale bar, the north arrow, the projection note, the classification marking, and the product identifier. The legend is the consumer's key to the map, and a legend that is incomplete, unclear, or misplaced undermines the map's usability.

The map element layout places the title at the top, the legend in the lower right, the scale bar and north arrow near the legend, the classification marking at the top and bottom of every page, and the

product identifier and the production date in the footer. The layout is specified in the cartographic template and applied automatically by the product studio.

The United States Geological Survey maintains cartographic standards for its published maps through the USGS National Geospatial Program's Standards for Digital Orthoimages, Digital Elevation Models, and Hydrography. The Ordnance Survey maintains cartographic standards through the Ordnance Survey Map Specification, which specifies the symbology (the OS MasterMap topographic layer's feature codes and their corresponding symbols), the generalization (the OS MasterMap's generalization rules for each scale), the legend layout, and the map element layout. The French IGN maintains cartographic standards through the IGN's Cartographic Specifications, which specify the same elements for the IGN's map series. The geospatial office maintains its cartographic standards, derived from the government's standards where such standards exist and from the national mapping agencies' standards where the government's standards do not exist.

11.3 Symbology and Classification Marking

Symbology is the set of symbols, colours, and patterns used to represent the features on the map. It is the visual language of the office's intelligence. The symbology must be consistent across products of the same type, so that the consumer can interpret the map without relearning the symbology for each product. The symbology must be distinguishable, so that the consumer can differentiate the features represented by different symbols. The symbology must be appropriate to the feature, so that the consumer can intuit the feature from the symbol.

The symbology is specified in the cartographic standards and applied by the production operator through the product studio's symbology library. The symbology library is a collection of symbol definitions (point symbols, line symbols, fill symbols, and label styles) that are applied to the features based on the feature's attributes. The library is maintained by the production operator, with the technical lead reviewing the symbol definitions and the office manager approving the library for production use. When a new product type is introduced, the production operator develops the symbology for the new type, reviews it with the technical lead and the specialist, and adds it to the library after the office manager's approval.

The symbology for the office's principal products is as follows. The flood inundation map uses the blue depth gradient (described in section 11.2), with the parcel boundaries overlaid in black and the affected infrastructure highlighted in red. The crop condition bulletin uses the green-to-red condition gradient, with the parcel boundaries overlaid in grey and the district and regional boundaries in black. The infrastructure risk map uses the red encroachment symbols (triangles for vegetation encroachment, squares for third-party construction, circles for ground subsidence), with the infrastructure assets in black and the vulnerability scores as labels. The prospectivity map uses the warm-cool prospectivity gradient, with the geological formations overlaid in muted browns and greens and the known deposits marked as stars.

Classification marking is the process by which the product's classification level is marked on the product. Every product carries a classification marking, and the marking is determined by the sensitivity of the product's content. The marking is applied by the production operator, reviewed by the office manager, and enforced by the catalog and archive submodule. No product is released without the correct marking.

The classification marking is applied in the standard format and the standard location. The format is the classification level in capital letters (PUBLIC, INTERNAL, CONFIDENTIAL, MEDICAL CONFIDENTIAL, SECURITY CONFIDENTIAL), centred at the top and bottom of every page of the product, on the cover, and in the metadata. For commercial products, the CONFIDENTIAL marking includes the customer identity and the product identifier. For Security Confidential products, the marking includes the handling caveats (e.g., "Not releasable to foreign nationals"). The marking in the metadata is embedded in the GeoTIFF's metadata tags or the PDF's document properties, so that the classification is machine-readable and can be verified by the consumer's information system.

The classification marking is particularly important for the energy infrastructure and minerals domains. The energy infrastructure products carry the SECURITY CONFIDENTIAL marking by default, because the geolocation of critical energy infrastructure is security-sensitive. The minerals products carry the CONFIDENTIAL marking for commercial commissions, because the customer's identity and the specific findings are commercially sensitive. The office manager reviews the classification marking on every product before release, and the review is the quality gate that prevents the release of misclassified products. A product that is over-classified (marked SECURITY CONFIDENTIAL when it should be PUBLIC) restricts the intelligence from consumers who should have access. A product that is under-classified (marked PUBLIC when it should be SECURITY CONFIDENTIAL) discloses sensitive intelligence to consumers who should not have access. Both errors are serious, and both are prevented by the office manager's review.

11.4 Format Conversion

Format conversion is the process by which the product is converted from the production format to the formats required by the consumers. The consumers have different format requirements, and the office must produce the product in the format that the consumer's system can use. A product that is produced in the wrong format, or in a format that the consumer's system cannot read, will sit unused on the consumer's desk; the analytical work that went into the product is wasted, and the consumer's decision is made without the office's intelligence.

The principal geospatial formats the office produces are GeoTIFF for raster products, Shapefile and GeoPackage for vector products, and KML/KMZ for products that the consumer will view in a virtual globe (such as Google Earth or Cesium). GeoTIFF is the standard format for raster geospatial data: it embeds the georeferencing information (the coordinate reference system, the pixel size, the extent) in the TIFF file's metadata tags, so that the consumer's GIS software can read the georeferencing automatically. GeoTIFF is supported by all major GIS software (ArcGIS, QGIS, GRASS GIS, Google Earth Engine). Shapefile is the legacy vector format developed by Esri for ArcGIS; it is widely supported but has limitations (the 2-gigabyte file size limit, the 10-character field name limit, the use of multiple files with the same base name). GeoPackage is the open standard vector format (OGC Standard 12-128r14) that addresses the Shapefile's limitations; it is a single-file SQLite database that can store vector features, raster tiles, and attributes. KML (Keyhole Markup Language) is the OGC standard for virtual globe display; it embeds the geometry, the symbology, and the attributes in an XML format that Google Earth and Cesium can read.

The principal report formats are PDF for printed reports and DOCX for reports that the consumer will edit. PDF (Portable Document Format, ISO 32000-2) preserves the layout, the fonts, and the graphics across platforms and is the standard format for distributing read-only documents. DOCX (the Microsoft Word Open XML format, ECMA-376) allows the consumer to edit the report, which may be necessary if the consumer needs to extract data from the report or to incorporate the report's content into the consumer's own documents.

The format conversion preserves three properties across all formats. The georeferencing is preserved for the geospatial formats: the GeoTIFF's metadata tags, the Shapefile's projection file (.prj), the GeoPackage's spatial reference system, and the KML's coordinate system are all correct and consistent with the office's standard projection. The symbology is preserved for the report formats: the PDF's colours, line weights, and fonts match the cartographic template, and the DOCX's styles are applied so that the report is visually consistent with the PDF. The classification marking is preserved for all formats: the marking is applied in the PDF's header and footer, in the DOCX's header and footer, in the GeoTIFF's metadata, in the Shapefile's metadata, in the GeoPackage's metadata, and in the KML's description. A GeoTIFF without its georeferencing is unusable; a PDF without its symbology is unreadable; a product of any format without its classification marking is a security breach.

The format conversion is conducted through the product studio, which provides the format conversion tools that the production operator uses to convert the product from the production format to the consumer-required formats. The tools are maintained by the production operator, with the technical lead reviewing the tool configuration and the office manager approving the tools for production use. The format conversion is the last step in the production function before the product is submitted to the office manager for review.

The United States Geological Survey provides its products in multiple formats through the EarthExplorer platform and the USGS data portals: GeoTIFF for Landsat and elevation data, Shapefile and GeoPackage for the National Hydrography Dataset and the National Land Cover Database, and PDF for the topographic maps. The European Space Agency's Copernicus programme provides its Sentinel products in the Sentinel-SAFE format (the ESA standard for Sentinel data) and in GeoTIFF through the Copernicus Open Access Hub. The Ordnance Survey provides its products in multiple formats through the OS Data Hub: GML (Geography Markup Language) and GeoPackage for the vector products, GeoTIFF for the raster products. The geospatial office provides its products in the formats required by its consumers, with the format conversion conducted through the product studio.

11.5 The Publication Pipeline

The publication pipeline is the sequence of steps by which a model output is transformed into a released product. The pipeline comprises seven steps: the receipt of the model output from the modeling workbench, the application of the cartographic template, the application of the symbology, the application of the classification marking, the format conversion, the review by the office manager, and the release to the consumer. The pipeline is the sequence that every product passes through, and the sequence is the operational discipline that ensures every product is produced to the same standard.

Step 1: receipt of the model output. The production operator receives the model output from the modeling workbench. The operator verifies that the model output includes the analytical result (the inundation map, the crop condition assessment, the change detection result, or the prospectivity map), the validation record, and the specialist's authorization for production. If the model output is incomplete — if the validation record is missing or the specialist has not authorized the production — the operator returns the model output to the specialist and does not proceed with the production.

Step 2: application of the cartographic template. The operator selects the cartographic template that matches the product type (flood inundation map, crop condition bulletin, infrastructure risk map, prospectivity study) and the classification level. The template specifies the map projection, the page layout, the legend layout, the scale bar and north arrow placement, the title block, and the classification marking locations. The operator applies the template to the model output through the product studio.

Step 3: application of the symbology. The operator applies the symbology from the symbology library to the features in the model output. The symbology is applied automatically by the product studio based on the feature attributes, but the operator verifies that the symbology is correct by visually inspecting the map: the depth gradient is applied to the flood inundation, the condition gradient is applied to the crop condition, the change symbols are applied to the detected changes, the prospectivity gradient is applied to the prospectivity scores.

Step 4: application of the classification marking. The operator applies the classification marking in the standard format and the standard location. The marking is determined by the product's content and the domain's classification regime: the flood inundation map may be Public (for the general flood hazard) or Internal (for the infrastructure exposure) or Security Confidential (for the dam failure scenario); the crop condition bulletin is Public (for the summary) or Internal (for the detailed assessment); the infrastructure risk map is Security Confidential; the prospectivity map is Public (for the base geology) or Confidential (for the commercial commission) or Security Confidential (for the critical minerals inventory).

Step 5: format conversion. The operator converts the product from the production format to the formats required by the consumer. The consumer's format requirements are specified in the product request or in the standing arrangement between the office and the consumer. The emergency management authority may require GeoTIFF (for ingestion into its operational mapping system) and PDF (for the evacuation guidance). The strategic reserve authority may require PDF (for the crop condition bulletin) and DOCX (for the procurement advisory, which the authority may incorporate into its own briefing). The energy sector entity may require GeoTIFF and Shapefile (for ingestion into its asset management system). The commercial minerals customer may require GeoTIFF (for the prospectivity map) and PDF (for the commission report).

Step 6: review by the office manager. The office manager reviews the product for compliance with the cartographic standards (is the symbology correct? is the generalization appropriate? is the legend complete?), for correctness of the classification marking (is the marking correct for the product's content? is the marking applied in the standard format and location? is the marking consistent across all components?), for usability of the format (does the GeoTIFF have its georeferencing? does the PDF have its symbology? does the product have its classification marking?), and for consistency with the office's other products (does the product use the same

symbology as previous products of the same type? does the product use the same layout?). The product that passes the review is authorized for release; the product that fails is returned to the production operator for correction.

Step 7: release to the consumer. The operator releases the product to the consumer through the appropriate distribution channel. For public-service products, the operator releases through the public-service distribution channel (the secure data interface, the web portal, or the direct delivery to the consumer's system). For commercial products, the operator releases through the commercial distribution channel, but only after the office manager has signed the commercial release log under the commercial release procedure (section 13.4). The release is recorded in the catalog and archive submodule, with the product identifier, the release date, the consumer, the distribution channel, and the office manager's authorization.

11.6 The Cartographic Standards

The cartographic standards specify the symbology, the generalization, the legend, and the layout for the principal map types the office produces. The standards are maintained by the production operator, with the technical lead reviewing and the office manager approving. The standards are derived from the government's cartographic standards, where such standards exist, and from the comparative practice of national mapping agencies, where the government's standards do not exist.

The standards specify four elements for each map type. The symbology specification defines the symbols, colours, and patterns for each feature type, with the colour values specified in the CMYK colour model (for print products) and the RGB colour model (for digital products). The generalization specification defines the generalization rules for each map scale, with the rules implemented as automated generalization tools in the product studio. The legend specification defines the legend layout, the legend entries (in the order of the features' importance), and the legend's location on the map. The layout specification defines the title block, the scale bar, the north arrow, the projection note, the classification marking locations, the product identifier, and the production date.

The cartographic standards are not static; they are updated as the office's products evolve and as the cartographic practice advances. The updates are proposed by the production operator, reviewed by the technical lead, and approved by the office manager, and the updates are recorded in the office's procedure reference. The procedure reference is the office's controlled document that records all changes to the standards, with the change date, the change description, the proposer, the reviewer, and the approver. The procedure reference is the audit trail that the national audit authority can examine to verify that the standards are maintained through a controlled process.

The United States Geological Survey maintains its cartographic standards through the National Geospatial Program's Standards documents, which are updated through the Survey's standards review process (a public review process that includes federal, state, local, and academic stakeholders). The Ordnance Survey maintains its cartographic standards through the OS Map Specification, which is updated through the Ordnance Survey's product management process. The French IGN maintains its cartographic standards through the IGN's Cartographic Specifications, which are updated through the agency's standards committee. The geospatial office maintains its cartographic standards and updates them through its standards review process, with the updates recorded in the office's procedure reference. The standards review process is initiated when a new product type is introduced, when a

consumer requests a change to the symbology or the layout, when a change in the cartographic practice (such as a new colour standard or a new generalization algorithm) warrants an update, or when the annual procedure review (section 32.3) identifies a need for revision. The review process ensures that the standards remain current with the office's products, the consumers' needs, and the cartographic practice.

11.7 Production and the Consumer

The production function produces the products that the consumer uses. The product is the office's intelligence delivered to the consumer, and the consumer's use of the product is the test of the production function. A product that is analytically correct but produced poorly — with poor symbology, with the wrong format, with the wrong classification marking — will not be used, and the analytical work that went into the product is wasted.

The consumer's decision is the test of the product. The emergency management authority's evacuation order is based on the flood inundation map. The strategic reserve authority's procurement decision is based on the crop condition bulletin and the procurement advisory. The energy sector entity's maintenance dispatch is based on the infrastructure risk map. The commercial minerals customer's exploration decision is based on the prospectivity map. Each decision is made on the basis of the product as delivered, and the product as delivered is the production function's responsibility.

The production function is the last capability in the pipeline before the archive. The acquisition capability acquires the data; the base layer capability maintains the geographic context; the processing capability transforms the data; the modeling capability produces the analysis; the production capability produces the product. Each capability depends on the capabilities before it, and the production capability depends on all of them. A failure in any earlier capability appears in the production capability as a product that cannot be produced to standard, and the production operator is often the first to detect the failure, because the operator is the first to see the model output in its final form.

The production function is the responsibility of the production operator, who operates the production function on a daily basis, with the office manager reviewing and the technical lead configuring. The operator's discipline is the foundation of the production function's quality, maintained through the training pipeline, the qualifications framework, and the professional identity treated in Part I. The production discipline is the last discipline of the pipeline before the archive, and the office's intelligence delivery depends on it. The operator who maintains the discipline produces products that are clear, correctly classified, and in the right format; the operator who does not produces products that are unclear, misclassified, or in the wrong format, and the products sit unused on the consumer's desk while the consumer's decision is made without the office's intelligence.

11.8 Worked Example: Producing a Flood Inundation Map for Event Response

The hazard specialist has completed the flood model run (section 10.8) and has authorized the model output for production. The production operator receives the model output from the modeling workbench: the inundation extent (a polygon), the flood depth (a 5-metre raster), the exposure report

(347 affected parcels, 1,200 people, 2 substations, 1 transmission corridor), and the validation record (peak stage within 0.13 metres at 3 gauges, inundation extent 87 per cent true positive rate).

The operator selects the cartographic template for the flood inundation map with Internal classification (because the exposure report includes critical infrastructure — the substations and the transmission corridor). The template specifies the UTM Zone 42N projection, the A3 landscape page layout, the legend in the lower right corner, the scale bar and north arrow adjacent to the legend, the title block at the top ("Flood Inundation Assessment — Riverbourne Basin — 17 January 2024"), and the classification marking at the top and bottom of every page.

The operator applies the symbology. The flood depth raster is symbolized with the blue depth gradient: 0 to 0.5 metres in light blue (CMYK 40/0/0/0), 0.5 to 1.5 metres in medium blue (CMYK 60/0/0/0), over 1.5 metres in dark blue (CMYK 80/0/0/20). The parcel boundaries are overlaid in black (CMYK 0/0/0/100) at 0.5-point line weight. The affected parcels are highlighted with a red fill (CMYK 0/100/100/0) at 50 per cent transparency. The substations are marked with red squares (8-point), and the transmission corridor is marked with a red line (2-point). The evacuation routes (from the emergency management authority's road network) are overlaid in green (CMYK 100/0/100/0) at 1.5-point line weight.

The operator applies the classification marking. The marking is "INTERNAL" in capital letters, centred at the top and bottom of every page. The marking includes the product identifier (PUB-2024-119-FLOOD) and the production date (17 January 2024). The marking is also embedded in the GeoTIFF's metadata tags and the PDF's document properties.

The operator conducts the format conversion. The consumer (the emergency management authority) has requested the product in three formats: GeoTIFF (for ingestion into the authority's operational mapping system, which uses Esri ArcGIS), Shapefile (for the affected parcels, which the authority's GIS team will use for the parcel-level evacuation planning), and PDF (for the evacuation guidance, which the authority will distribute to the field teams). The operator converts the flood depth raster to GeoTIFF, the affected parcels polygon to Shapefile (with the .prj projection file, the .shp, the .shx, and the .dbf files), and the full product (the map layout with the symbology, the legend, and the classification marking) to PDF.

The operator verifies the format conversion: the GeoTIFF's georeferencing is correct (UTM Zone 42N, 5-metre pixel size, the extent matches the model output), the Shapefile's projection file specifies UTM Zone 42N (EPSG:32642), and the PDF's classification marking is applied at the top and bottom of every page. The operator submits the product to the office manager for review.

The office manager reviews the product. The manager checks the cartographic compliance: the symbology matches the cartographic standards (the blue depth gradient, the black parcel boundaries, the red infrastructure highlights), the generalization is appropriate for the 1:25,000 scale (building footprints and road centrelines are preserved), the legend is complete (the depth gradient, the parcel boundaries, the infrastructure highlights, and the evacuation routes are all listed), and the layout matches the template (the title, the legend, the scale bar, and the north arrow are in their standard locations). The manager checks the classification marking: the "INTERNAL" marking is applied at the top and bottom of every page, the product identifier and the production date are included, and the marking is consistent across the GeoTIFF, the Shapefile, and the PDF. The manager checks the

format usability: the GeoTIFF has its georeferencing, the Shapefile has its projection file, and the PDF has its symbology. The manager authorizes the release.

The operator releases the product to the emergency management authority through the secure data interface. The release is recorded in the catalog and archive submodule: the product identifier (PUB-2024-119-FLOOD), the release date (17 January 2024, 14:30 UTC), the consumer (Emergency Management Authority), the distribution channel (secure data interface), the formats (GeoTIFF, Shapefile, PDF), the classification marking (INTERNAL), and the office manager's authorization. The operator archives a copy of the product in all three formats, with the provenance record linking the product to the model run (HEC-RAS 2D 6.3.1, the event parameters, the validation results), the processing run (the LiDAR DEM, the land cover data, the hydrography data), and the acquisition (the sensor feed data, the Sentinel-1 SAR image used for the validation). The chain is unbroken from the sensor feeds through the hydrological model, the hydraulic model, the inundation map, the production run, and the archive, to the consumer's operational system.

Chapter 12 — Catalog and Archive

12.1 The Archive Function

Catalog and archive is the sixth and final capability of the pipeline. It preserves the office's outputs and maintains the catalog that makes them retrievable. The capability maintains the complete provenance chain from acquisition through archive, with version history, classification markings, retention designations, and access logging for classified products. The archive is the institutional memory of the office, and it is the record that allows the office to reconstruct its past work, audit its current work, and plan its future work. The archive is subject to the audit jurisdiction of the national audit authority, and the audit is the external verification that the archive is maintained.

The archive function is conducted through the catalog and archive submodule of the geospatial information system, treated in Chapter 29. The submodule manages the product catalog, the version history, the classification markings, the retention designations, and the access logging. It provides the interface through which the production operator archives the products, the office manager ensures compliance, and the specialists retrieve the products for re-use or re-analysis. The design and use of the catalog and archive submodule are specified in the job breakdown sheets in Chapter 16.

The archive function is operated by the production operator on a daily cycle. The operator receives the released products from the production function, records them in the catalog with their product identifier, type, production date, classification marking, retention designation, provenance record, and access permissions, and stores them in the archive. The office manager ensures the archive's compliance with the records management standards and the classification framework, and the national audit authority verifies the compliance periodically. The specialists retrieve the archived products through the catalog and archive submodule for re-use in new analyses or for re-analysis with updated methods.

The archive function has five sub-capabilities, each treated in a separate section of this chapter. The catalog (section 12.2) is the index of the office's products. Version history (section 12.3) is the record of the versions of the office's products. Retention and disposal (section 12.4) manages the lifecycle of the archived products. Access control (section 12.5) manages the access to the archived products. The audit jurisdiction (section 12.6) is the external verification that the archive function is maintained.

12.2 The Catalog

The catalog is the index of the office's products. It comprises the product identifier, the product type, the production date, the classification marking, the retention designation, the provenance record, and the access permissions for every product the office has ever produced. The catalog is maintained by the production operator, with the technical lead configuring the catalog structure and the office manager ensuring compliance. The catalog allows the office to answer the question "what products have we produced?" for any period of its history, and to retrieve the products that match the query.

The catalog is structured as a database with the product identifier as the primary key and the product type, the production date, the classification marking, the domain, and the retention

designation as the secondary keys. The primary key is a unique identifier assigned to each product at the moment of archiving: the format is a prefix (indicating the domain: FLOOD, CROP, INFR, MINE), a year, a sequence number, and a version number, separated by hyphens (e.g., FLOOD-2024-017-v01). The secondary keys allow the catalog to be queried by product type (retrieve all flood inundation maps), by date range (retrieve all products produced in January 2024), by classification (retrieve all Security Confidential products), by domain (retrieve all minerals products), or by retention designation (retrieve all products with permanent retention).

The catalog's query capabilities support the office's operational needs and its audit obligations. The operational queries include the retrieval of products for re-use (a specialist queries the catalog for all flood inundation maps of the Riverbourne Basin to calibrate a new hydrological model), the retrieval of products for re-analysis (a specialist queries the catalog for all crop condition bulletins from the 2023 growing season to re-analyze them with an updated crop model), the retrieval of products for consumer inquiries (the office manager queries the catalog for the infrastructure risk map produced in October 2023, which the energy sector entity has requested), and the retrieval of products for the office's annual reporting (the office manager queries the catalog for all products produced in the calendar year, with their classification markings and their distribution records, for the annual report to the desk chief). The audit queries include the retrieval of products for the national audit authority's audit (the auditor queries the catalog for all commercial minerals products produced in the audit period, with their provenance records and their commercial release log entries) and the retrieval of products for the office's internal audit (the office manager queries the catalog for all products that failed the validation and were reprocessed or rejected, to verify that the quality assurance process is functioning).

The United States Geological Survey maintains a catalog of its published products through the Survey's data management system, and the catalog is the mechanism that makes the Survey's products retrievable through the EarthExplorer platform and the National Map viewers. The catalog uses the Global Visualization Viewer (GloVis) for browsing and the EarthExplorer for searching, with the products indexed by dataset (Landsat, Sentinel, MODIS), by acquisition date, by path/row (for Landsat), by tile (for Sentinel), and by cloud cover. The European Space Agency's Copernicus programme maintains a catalog of its products through the Copernicus Open Access Hub, with the products indexed by mission (Sentinel-1, Sentinel-2, Sentinel-3), by acquisition date, by orbit, by tile, and by processing level. The Ordnance Survey maintains a catalog through the OS Data Hub, with the products indexed by product type (OS MasterMap, OS OpenData, OS Tiles), by tile, and by update date. The geospatial office maintains its catalog through the catalog and archive submodule, with the catalog structured to support both the operational needs and the audit obligations.

12.3 Version History

Version history is the record of the versions of the office's products. It comprises the version number, the version date, the version description, and the version provenance for every version of every product. The version history is maintained by the production operator, with the technical lead configuring the version management and the office manager ensuring compliance. The version history allows the office to answer the question "what version of this product did we produce, and

when?" for any product it has ever produced, and to retrieve the specific version that was available at a specific point in time.

The office's products change over time, and the changes are the record of the office's evolving understanding of its domains. The flood inundation map is updated as the model is refined (a new Manning's n calibration improves the depth accuracy) and as new events are observed (a new flood event provides additional validation data). The crop condition bulletin is updated as the season progresses (each weekly bulletin supersedes the previous one with the latest imagery and the updated yield forecast) and as the harvest outcomes become available (the post-harvest bulletin incorporates the reported harvest and the model's accuracy assessment). The infrastructure risk map is updated as the monitoring cycle continues (each monthly cycle supersedes the previous one with the latest change detections) and as the ground truth inspections are conducted (the inspection results may correct the change detections). The mineral prospectivity map is updated as the geological survey advances (new geological mapping, new geophysical surveys) and as new exploration data is acquired (drill holes, trench samples).

The version numbering convention uses a major.minor scheme. A major version increment (v01 to v02) indicates a substantial revision: a new model version, a new data source, or a new analytical method. A minor version increment (v01 to v01.1) indicates a correction: a classification marking correction, a cartographic error fix, or a format conversion error fix. The version description summarizes the change ("v02: recalibrated Manning's n values based on the 2024 flood event validation; inundation extent true positive rate improved from 82 per cent to 87 per cent"). The version provenance records the model version, the parameters, the data inputs, and the validation results that produced the version, linking the version to the processing provenance and the acquisition provenance.

The version history is particularly important for the audit jurisdiction of the national audit authority, because the audit may require the reconstruction of the product that was available at a specific point in time. An audit that asks "what did the office know about the flood risk in district X on date Y?" is answered by retrieving the version of the flood inundation map that was current on date Y, with the version's provenance record showing the model version, the parameters, the data inputs, and the validation results that produced that version. The retrieval is conducted through the catalog, which maintains the version history for every product and allows the auditor to query the catalog by product identifier and date to retrieve the specific version.

The version history also supports the office's analytical improvement. The specialist who produces a new version of a product can compare the new version with the previous version to verify that the new version is an improvement: the new flood inundation map should have a higher true positive rate than the previous version; the new yield forecast should have a lower error relative to the reported harvest; the new change detection should have a higher detection rate. The comparison is the office's mechanism for verifying that its analytical methods are improving over time, and the comparison is only possible because the version history preserves the previous versions for comparison. Without the version history, the office would have no baseline against which to measure improvement, and the office would have no evidence to present to the national audit authority or to the office's consumers that the office's analytical methods are getting better. The version history is the

record of the office's analytical progress, and the record is the evidence that the office is learning from its experience.

12.4 Retention and Disposal

Retention and disposal is the process by which the geospatial office manages the lifecycle of its archived products, from the moment of archiving to the moment of disposal. The retention designation specifies the period for which the product is retained, and the disposal specifies the method by which the product is disposed of at the end of the retention period. The retention and disposal are governed by the government's records management standards, which specify the retention periods and the disposal methods for the principal product types. The records management standards are established by the national archives or the government's records management authority, and they apply to all government records, including the geospatial office's products. The standards are typically documented in a records retention schedule, which lists each record type, its retention period, and its disposal method. The geospatial office's products are classified in the retention schedule as scientific and technical records, with the retention periods and disposal methods specified by the schedule. The office manager is responsible for ensuring that the office's retention and disposal practices comply with the schedule, and the national audit authority verifies the compliance during the annual audit.

The retention periods vary by product type and by classification. The flood inundation maps are retained permanently, because they are the record of the office's flood intelligence and the basis for future flood risk assessment. A flood inundation map produced for the 2024 event will be used in 2034 to calibrate the hydrological model for the next event, and in 2044 to assess whether the flood risk has changed due to climate change or land use change. The crop condition bulletins are retained for ten years, because they are the record of the office's food security intelligence and the basis for future yield forecasting. A crop condition bulletin produced for the 2024 growing season will be used in 2034 to establish the historical baseline for the 2034 season's condition assessment. The energy infrastructure risk maps are retained for seven years, the period of the sectoral regulator's compliance cycle, because they are the record of the office's energy infrastructure intelligence and the basis for the regulator's compliance verification. The mineral prospectivity maps are retained permanently, because they are the record of the office's mineral intelligence and the basis for future prospectivity assessment. The commercial minerals commissions are retained for the period specified in the commission contract (typically ten years after the commission's completion), to support the customer's obligations and any contractual disputes.

The disposal is conducted through secure deletion. The deletion method varies by classification. For Public and Internal products, the deletion is a standard file deletion (the product's files are deleted from the archive's storage, and the catalog entry is marked as "disposed"). For Confidential products, the deletion is a secure file deletion (the product's files are overwritten with random data before deletion, to prevent recovery). For Security Confidential products, the deletion is a cryptographic erasure (the encryption key used to encrypt the product is destroyed, rendering the encrypted product permanently unreadable) or a physical destruction (if the product is stored on physical media, the media is destroyed in accordance with the classification framework's physical destruction standards).

The deletion is recorded in the catalog and archive submodule, with the disposal date, the disposal method, the witness (for classified products), and the authorizing officer.

The retention and disposal are the lifecycle management of the office's products, governed by the records management standards and enforced by the office manager. The disposal record is the evidence that the product was disposed of in accordance with the standards, and the evidence is the operational practice that establishes the office's compliance with the records management discipline. The national audit authority verifies the retention and disposal compliance during the annual audit, and the audit's findings may trigger corrective action where products have been retained beyond their retention period (a storage cost and a potential security risk) or disposed of before their retention period expired (a loss of institutional memory and a potential audit failure).

12.5 Access Control

Access control is the process by which the geospatial office manages the access to its archived products. It comprises the access permissions, the access logging, and the access review. The access permissions specify who may access each product, and the permissions are determined by the product's classification marking. The access logging records every access to the classified products, and the logging allows the office to answer the question "who accessed this product?" for any classified product. The access review is the periodic review of the access permissions and the access logs, and the review ensures the access control is maintained.

The access permissions are implemented through role-based access control (RBAC). Each officer in the office is assigned a role (office manager, technical lead, data pipeline operator, production operator, hazard specialist, resource economics specialist), and each role has a set of permissions for each classification level. The office manager has access to all products (Public, Internal, Confidential, Security Confidential). The technical lead has access to all products except the commercial minerals commissions (Confidential, restricted to the office manager and the resource economics specialist). The data pipeline operator has access to Public and Internal products. The production operator has access to all products (because the operator archives all products), but the operator's access to the Confidential and Security Confidential products is logged and reviewed. The hazard specialist has access to Public, Internal, and Security Confidential products (the Security Confidential energy infrastructure products). The resource economics specialist has access to Public, Internal, Confidential (the commercial minerals commissions), and Security Confidential products (the critical minerals inventory).

The access logging records every access to the Confidential and Security Confidential products. The log entry includes the access timestamp, the accessor's identity (the officer's user ID), the product identifier, the access type (read, download, modify), and the access purpose (re-use, re-analysis, audit, training). The access log is maintained in the catalog and archive submodule and is reviewed by the production operator daily (section 16.6) and by the office manager monthly. The daily review identifies any unauthorized access (access by an officer who does not have the required permissions) or any unusual access patterns (access outside working hours, bulk downloads, access by an officer who does not normally access the product type). The monthly review verifies that the access permissions are current (officers who have left the office have had their permissions revoked; officers

who have changed roles have had their permissions updated) and that the access log is complete (no gaps in the log that would indicate a logging failure).

The access control is particularly important for the Security Confidential and Confidential products. The Security Confidential products, including the energy infrastructure risk maps and the critical minerals inventory, are restricted to the security-cleared personnel, and the access is logged and reviewed. The Confidential products, including the commercial minerals commissions, are restricted to the office manager and the resource economics specialist, and the access is logged and reviewed. The access permissions are reviewed when personnel join the office (the new officer's permissions are configured based on the officer's role), when personnel leave the office (the departing officer's permissions are revoked), and when personnel change roles within the office (the officer's permissions are updated to reflect the new role).

The United States Geological Survey's access control for its classified products is governed by the federal government's classified information management framework (Executive Order 13526, "Classified National Security Information," and 32 CFR Part 2001, the National Industrial Security Program). The Ordnance Survey's access control for its commercial products is governed by the Ordnance Survey's commercial data management framework, which is specified in the PSMA and the commercial license terms. The national mapping agencies maintain their access control for their classified and commercial products through their data management frameworks. The geospatial office maintains its access control through the catalog and archive submodule, with the access permissions, the access logging, and the access review configured for each product's classification.

12.6 The Audit Jurisdiction

The audit jurisdiction of the national audit authority is the external verification that the archive function is maintained. The national audit authority examines the archive's completeness, the version history's accuracy, the retention and disposal's compliance, and the access control's effectiveness. The examination is conducted periodically (annually for most national audit authorities), with the findings published in the audit authority's report. The audit's findings can trigger legislative or administrative action where the archive has been eroded.

The audit examines four questions. First, is the archive complete, with every product the office has produced recorded in the catalog? The auditor queries the catalog for all products produced in the audit period and compares the count with the production log's count. A discrepancy (the production log shows more products than the catalog) indicates that products were produced but not archived, which is a records management failure. Second, is the version history accurate, with every version of every product recorded with its version number, date, description, and provenance? The auditor retrieves a sample of products from the catalog and verifies that the version history is complete and that the provenance record links to the processing and acquisition provenance. Third, is the retention and disposal compliant with the records management standards, with the retention periods and the disposal methods correctly applied? The auditor retrieves a sample of products whose retention period has expired and verifies that they have been disposed of in accordance with the standards, with the disposal recorded in the catalog. Fourth, is the access control effective, with the access permissions correctly assigned, the access logging correctly maintained, and the access review correctly conducted? The auditor examines the access log for a sample of classified products, verifies

that all access was by authorized personnel, and checks the access review records for the daily and monthly reviews. The daily review is conducted by the production operator at the end of each working day and is specified in the job breakdown sheet in section 16.6. The monthly review is conducted by the office manager and includes a comprehensive examination of the access log for all classified products, a verification of the access permissions for all office personnel, and a documentation of any access anomalies and their resolution.

The Government Accountability Office (GAO), established in 1921, audits the United States Geological Survey's records management as part of its broader audit of the Department of the Interior's information management. The GAO's audit findings are published in GAO reports (e.g., GAO-15-680, which examined the Survey's enterprise activities and identified issues with the revolving fund management and the cost accounting for commercial activities). The National Audit Office (NAO), established in 1983, audits the Ordnance Survey's trading fund as part of its broader audit of central government trading funds. The NAO's audit findings are published in NAO reports and are reviewed by Parliament's Public Accounts Committee. The Cour des comptes, established in 1807, audits IGN France's records management as part of its broader audit of the government's information systems. The Australian National Audit Office (ANAO), established in 1997, audits Geoscience Australia's records management as part of its broader audit of the agency's financial management.

The audit's independence from the office is the condition of its credibility. The audit authority reports to the legislature, not to the executive, and the audit authority's budget is appropriated by the legislature, not by the executive. This institutional independence ensures that the audit authority cannot be pressured by the office's parent ministry to soften its findings or to omit issues that might embarrass the ministry. An office that maintains its archive welcomes the audit, because the audit's clean findings are the evidence the office needs to demonstrate its compliance to its consumers and to the public. An office that has eroded its archive dreads the audit, because the audit's adverse findings are the trigger for legislative or administrative action that may restructure the office, replace its leadership, or revoke its authorities.

12.7 The Archive and the Office

The archive is the institutional memory of the office. It allows the office to learn from its past work, to audit its current work, and to plan its future work. The archive is the foundation of the office's institutional continuity, and the office's work depends on the archive's integrity. An office that loses its archive loses its memory: it cannot reconstruct the basis of its past decisions, it cannot audit the conduct of its past work, and it cannot plan its future work on the basis of its past experience. The office must reconstruct the lost work from the sources, where the sources are available, and the reconstruction is costly and incomplete.

The archive is used in three operational modes. In the re-use mode, the specialist retrieves a past product to use as the input to a new analysis: a past flood inundation map is retrieved to calibrate a new hydrological model; a past crop condition bulletin is retrieved to compare with the current season's condition; a past infrastructure risk map is retrieved to identify the changes that have occurred since the past map. In the audit mode, the office manager or the national audit authority retrieves a past product to verify the conduct of the office's past work: the audit verifies that the

product was produced in accordance with the standards, that the classification marking was correct, and that the access was controlled. In the training mode, the specialist uses the past products as case studies for the training of new specialists, and the past products are the office's accumulated experience made available to the new generation.

The archive is the last capability in the pipeline. The acquisition capability acquires the data; the base layer capability maintains the geographic context; the processing capability transforms the data; the modeling capability produces the analysis; the production capability produces the product; the archive capability preserves the record. The six capabilities are the pipeline, and the archive capability is the capability that preserves the record for the future. The preservation allows the office to build on its past work rather than repeat it: a flood inundation map produced for a past event is the input to the hydrological model's calibration for a future event; a crop condition bulletin produced for a past season is the input to the yield model's calibration for a future season.

The archive function is the responsibility of the production operator, who operates the archive function on a daily basis, with the office manager ensuring compliance and the technical lead configuring. The operator's discipline is the foundation of the archive function's integrity, maintained through the training pipeline, the qualifications framework, and the professional identity treated in Part I. The archive discipline is the last discipline of the pipeline, and the office's institutional memory depends on it. An office that maintains its archive maintains its memory: the office can reconstruct the basis of its past decisions, can audit the conduct of its past work, and can plan its future work on the basis of its past experience. An office that does not maintain its archive loses its memory: the office cannot reconstruct the basis of its past decisions, cannot audit the conduct of its past work, and must plan its future work without the benefit of its past experience, repeating the work that the archive should have preserved. The archive is the pipeline's final output, and the archive's integrity is the condition of the office's institutional continuity over time.

12.8 Worked Example: Archiving and Retrieving a Flood Inundation Map

The production operator has just released the flood inundation map for the Riverbourne Basin event (product identifier PUB-2024-119-FLOOD, released 17 January 2024 at 14:30 UTC through the secure data interface to the emergency management authority). The operator now archives the product.

The operator archives a copy of the product in all three produced formats (GeoTIFF, Shapefile, PDF) in the catalog and archive submodule. The operator records the catalog entry: product identifier (FLOOD-2024-017-v01), product type (flood inundation map), production date (17 January 2024), classification marking (INTERNAL), domain (flood risk), retention designation (permanent), access permissions (all office staff), provenance record (linked to the model run: HEC-RAS 2D 6.3.1, event parameters, validation results; linked to the processing run: LiDAR DEM 2m, land cover 30m, hydrography NHD; linked to the acquisition: sensor feeds from 15 rain gauges and 8 river stage sensors, Sentinel-1 SAR image S1B_IW_SLC__1SDV_20240117T053000). The catalog entry is the primary record that allows the product to be retrieved.

Six months later, the hazard specialist is calibrating the hydrological model for the next flood season. The specialist queries the catalog for all flood inundation maps of the Riverbourne Basin. The catalog returns two products: FLOOD-2024-017-v01 (the January 2024 event) and FLOOD-2023-

003-v02 (a March 2023 event, with the v02 indicating a post-event recalibration). The specialist retrieves both products, with their provenance records, and uses them to calibrate the HEC-HMS model's SCS curve number and the HEC-RAS 2D model's Manning's n values. The calibration is the re-use mode: the past products are the input to the new analysis.

One year later, the national audit authority conducts its annual audit. The auditor queries the catalog for all products produced in the calendar year 2024, with their provenance records and their distribution records. The catalog returns 127 products: 43 flood risk products, 52 food security products, 19 energy infrastructure products, and 13 minerals products. The auditor examines a sample of 15 products (selected to include all four domains and all four classification levels). For each sampled product, the auditor verifies: (a) the catalog entry is complete (product identifier, type, date, classification, retention, provenance, access permissions); (b) the provenance record links to the processing and acquisition provenance (the chain is unbroken from the source data through the processing, modeling, and production to the archive); (c) the classification marking is correct for the product's content; (d) the access log shows only authorized access; (e) the retention designation is correct for the product type. The auditor finds 14 of 15 products fully compliant. The 15th product (a crop condition bulletin) has a minor provenance gap: the processing provenance does not record the Sen2Cor processor version. The auditor notes the finding, and the office manager directs the technical lead to verify the Sen2Cor version from the processing log and add it to the provenance record. The correction is documented, and the audit's finding is resolved.

Two years later, the emergency management authority requests the flood inundation map that was available on 17 January 2024, for a post-event review of the evacuation decisions. The office manager queries the catalog for the product identifier FLOOD-2024-017, with the date 17 January 2024. The catalog returns FLOOD-2024-017-v01, which is the version that was current on that date (no subsequent versions have been produced, because the event was a one-time occurrence and the map has not been updated). The office manager retrieves the product, verifies the classification marking (INTERNAL), and delivers it to the emergency management authority through the secure data interface. The retrieval is the audit mode: the past product is retrieved to verify the conduct of the office's past work and to support the consumer's post-event review.

This worked example illustrates the archive function's five sub-capabilities operating together. The catalog (section 12.2) indexed the product and allowed it to be retrieved by identifier, by date, and by domain. The version history (section 12.3) recorded the version number and the provenance, allowing the specific version to be retrieved. The retention and disposal (section 12.4) retained the product permanently (flood inundation maps are permanent retention). The access control (section 12.5) restricted access to authorized office staff and logged the retrieval. The audit jurisdiction (section 12.6) verified the archive's completeness and the provenance chain's integrity, and the audit's finding (the missing Sen2Cor version) was corrected and resolved. The archive function preserved the product for re-use (the model calibration), for audit (the national audit authority's examination), and for the consumer's post-event review (the emergency management authority's request), and the preservation is the institutional memory that the pipeline is designed to establish. The archive function completes the pipeline: data is acquired, processed, modeled, produced, and then preserved, and the preservation ensures that the work of the pipeline endures beyond the moment of the product's release, serving the office and its consumers for years and decades to come.

Chapter 13 — The Office Manager

Part III — Roles and Daily Work: The six seats in the office, with Training Within Industry Job Breakdown Sheets for the major daily tasks. Each sheet specifies the important steps, the key points, and the reasons for the key points.

13.1 The Role of the Manager

The office manager heads the geospatial office. The manager is accountable for the office's production, its compliance with standards, its financial administration, and its maintenance of the commercial firewall. The position is established by the office's enabling instrument, and the manager is appointed by the bureau chief on the recommendation of the desk chief. The manager reports to the desk chief, who reports to the bureau chief. Performance is assessed by the desk chief against the standards set out in this Handbook.

The manager is the sole authority for commercial releases. The authority cannot be escalated upward to the desk chief or bureau chief, and it cannot be delegated downward to the technical lead or to a specialist. The single point of authority is the operational expression of the commercial firewall, treated in Chapter 5. The firewall is maintained through four mechanisms: separate charge codes for commercial and public-service work; separate classification markings for commercial and public-service products; the manager's sole authority for commercial release; and the commercial release log, which records every release and is examined by the national audit authority during the annual audit. The mechanisms are described in operational detail in the job breakdown sheets in sections 13.3 through 13.6.

The manager's accountability and the technical lead's accountability are separable but interdependent. The technical lead is accountable for the office's technical integrity: the pipeline configuration, the model validation, the systems integrity, and the specialist support. The manager is accountable for the office's institutional integrity: the production against standards, the firewall, the classification correctness, and the upward reporting. The two officers work in tandem. The manager authorizes releases on the basis of the technical lead's validation, and the technical lead's validation is conducted to standards the manager enforces. The two accountabilities are distinct, and confusion between them produces defective work.

The manager is the office's primary point of contact with the consumer offices. The emergency management authority, the strategic reserve authority, the energy sector entity, the sectoral regulator, and the commercial minerals customers address the manager for the office's products, escalations, and inquiries. The manager receives the consumer's request, directs it to the responsible specialist, and verifies the response before it is delivered. The manager is also the office's primary point of contact with the external authorities: the national audit authority for the annual audit, the cadastral and national mapping authorities for the base layer feeds, and the procurement authority for the framework agreements.

National mapping agencies organize the office-head function in comparable ways. The United States Geological Survey's center directors report to regional directors and exercise authority for center-level operations, including the center's publishing and partnership activities. Geoscience Australia's division heads report to the Chief Executive Officer and exercise authority for divisional

products, including the division's commercial geoscience services. IGN France's directions départementales report to regional directors and exercise authority for departmental products under the agency's commercial framework. The geospatial office's manager reports upward to the desk chief and exercises authority for the office's products, including the commercial minerals commissions, within the bureau's framework. The position is held by an officer who has completed the training pipeline, holds the senior certification, and has served as a technical lead or as a specialist for a minimum period specified by the qualifications framework. The position is a career position: the manager is appointed from within the role family, not from outside the geospatial profession. The career position ensures that the manager has the technical literacy to read the technical lead's validation, the operational literacy to read the production dashboard, and the institutional literacy to maintain the firewall. The appointment process for the manager position is specified in the office's enabling instrument. The bureau chief appoints the manager on the recommendation of the desk chief, following a competitive selection process that includes a review of the candidates' qualifications, certifications, experience, and references. The selection process is conducted in accordance with the government's civil service recruitment standards, and the appointment is documented in a formal instrument that specifies the manager's authorities, responsibilities, and reporting line. The manager's term is typically indefinite (a career appointment), but the manager may be removed for cause (misconduct, incapacity, or failure to meet the standards) through the government's disciplinary procedures. The removal process is documented in the government's civil service disciplinary regulations and requires the involvement of the civil service commission to ensure that the removal is justified and procedurally fair.

The manager's accountability is exercised through the office's procedures and verified through the audit jurisdiction. The procedures comprise the morning production review, the commercial release authorization, the monthly reporting cycle, and the incident response activation. Each procedure is a gate through which the office's work passes, and the manager's accountability is exercised at each gate. The audit jurisdiction is exercised by the national audit authority, which examines the office's archive, the commercial release log, the validation records, and the classification compliance during the annual audit. The audit's findings are reported to the bureau chief, and the findings may trigger administrative action where the office's work has fallen below the standards. The manager's accountability is also exercised through the desk chief's monthly review, which assesses the office's performance against the standards and addresses any unresolved incidents or exceptions.

13.2 The Manager's Daily Rhythm

The manager's day is organized around the office's two cadences: the fixed pipeline cadence and the variable modeling cadence. The fixed cadence comprises daily ingestion, the daily processing queue, the weekly base layer integrity check, the monthly archive audit, and production on demand. The variable cadence comprises the specialists' domain-specific model runs, which operate on the flood, food, energy, and minerals schedules. The two cadences run in parallel, and the manager oversees both.

The day begins with the morning production review (section 13.3). The review is conducted at the production dashboard with the data pipeline operator and the production operator present. The

review takes approximately fifteen minutes and covers the overnight production log, the satellite acquisitions received, the processing queue status, the modeling queue status, the production queue status, the flood event status, and the day's priorities. The review establishes the office's situational awareness for the day.

Commercial release authorizations (section 13.4) are conducted as commercial products become ready. A commercial release authorization is conducted for every commercial product, without exception, and the authorization is recorded in the commercial release log. The authorization verifies the firewall compliance, the classification marking, the charge code, the customer identity, and the agreed price before the product is released through the commercial channel.

Specialist support consultations occur as needed throughout the day. A specialist may request a consultation with the manager to discuss a consumer's request, a classification question, or a production timeline. The manager's role in these consultations is to provide institutional direction. The analytical work is the specialist's responsibility, and the technical lead supports the specialist's technical problems.

The day ends with a brief production review that verifies the day's work is complete and that any issues have been addressed or escalated. The review takes five to ten minutes. It confirms that the day's priorities were addressed, that the production queue is in the expected state, and that any incidents or exceptions are documented for the monthly reporting cycle (section 13.5).

The rhythm is adjusted to the office's operational needs. During a flood event, the event-driven cadence overrides the scheduled cadence, and the manager's day is organized around the incident response activation (section 13.6). During the monthly reporting cycle, the day is organized around the report compilation. During a quiet week, the day follows the standard rhythm. The adjustment is the manager's discretion, exercised in consultation with the technical lead.

The manager's weekly rhythm includes the weekly base layer integrity audit, conducted on the first working day of each week in coordination with the technical lead and the data pipeline operator. The audit verifies that the office's base layers are synchronized with the source authorities' current versions, and the manager's role is to review the audit's results and authorize any updates that affect the office's production. The manager's monthly rhythm includes the monthly archive audit, conducted on the first working day of each month, and the monthly reporting cycle. The archive audit verifies the archive's completeness and integrity, and the reporting cycle compiles the office's performance for the desk chief's review. The weekly and monthly rhythms complement the daily rhythm and ensure that the office's standards are maintained over time.

The manager's weekly rhythm includes a coordination meeting with the desk chief, conducted once a week, in which the manager provides a verbal update on the office's production, any incidents, and any resource needs. The coordination meeting is distinct from the monthly reporting cycle (section 13.5): the weekly meeting is informal and verbal, focused on immediate issues, while the monthly report is formal and documented, focused on performance against standards. The manager also conducts a weekly check-in with each specialist, lasting 10 to 15 minutes, in which the specialist reports on their current work, any technical issues, and any consumer requests that require the manager's attention. The weekly check-ins are the manager's mechanism for maintaining awareness of the specialists' work and for identifying issues before they become incidents. The daily rhythm is the manager's mechanism for maintaining situational awareness. A manager who skips the morning

production review lacks the overnight picture. A manager who skips the day-end production review does not confirm whether the day's priorities were addressed. The discipline is not formal: no checklist is filed, no log is signed. The discipline is the daily practice of being present at the dashboard and informed about the work.

13.3 The Morning Production Review

The morning production review is the manager's first task of the day. It is conducted at the start of the working day, before any other activity, and it takes approximately fifteen minutes. The review establishes the office's situational awareness for the day, identifies any issues that arose overnight, and sets the day's production priorities. The data pipeline operator and the production operator attend. The review is conducted at the production dashboard.

Important Steps	Key Points	Reasons for Key Points
1. Open the production dashboard and review the overnight production log.	Check the overnight ingestion and processing logs first, before reviewing the modeling or production logs.	Overnight failures in ingestion or processing cascade through the pipeline, and identifying them first prevents the office from working against a broken pipeline.
2. Verify that all scheduled satellite acquisitions were received.	Cross-check the acquisition schedule against the ingestion log; flag any missing acquisitions immediately.	Missing acquisitions break the pipeline for the affected domain, and the specialist needs to know before starting a model run on incomplete data.
3. Review the overnight processing queue for failed jobs.	Note the failure reason for each failed job; do not simply re-run without understanding the cause.	Re-running a failed job without understanding the failure reproduces the failure and wastes processing time.
4. Check the modeling queue for runs that failed validation.	Assign each failed run to the responsible specialist for investigation before the daily stand-up.	Failed validation runs indicate either a model problem or a data problem, and the specialist is the officer who can distinguish between them.
5. Review the production queue for products past their delivery deadline.	Flag any overdue products to the production operator and determine the cause of the delay.	Late products are a performance standard failure, and the cause must be identified and addressed before the product is released.

Important Steps	Key Points	Reasons for Key Points
6. Check the flood event status on the sensor feed dashboard.	Review the rain gauge and river stage alerts, and check the weather forecast for the next 48 hours.	Flood events can escalate quickly, and the office needs to pre-position its response before the emergency management authority calls.
7. Brief the office on the day's priorities.	State the top three priorities for the day clearly, and confirm that each priority has an assigned owner.	A six-person office needs shared situational awareness, and unassigned priorities are priorities that will not be addressed.

13.4 The Commercial Release Authorization

The commercial release authorization is the procedure by which the office manager authorizes the release of a commercial minerals product to a commercial customer. The procedure is conducted for every commercial product, without exception, and the authorization is recorded in the commercial release log. The manager is the sole authority. The authorization cannot be escalated upward to the desk chief or bureau chief, and it cannot be delegated downward to the technical lead or to a specialist.

Important Steps	Key Points	Reasons for Key Points
1. Receive the completed commercial product from the resource economics specialist.	Verify that the product is marked as complete and ready for release by the specialist.	An incomplete product must not enter the release procedure, and the specialist's completion mark is the signal that the analytical work is finished.
2. Review the product for firewall compliance.	Verify that the product was produced from commercial data sources and commercial model runs, not from public-service data or public-service model runs.	The firewall requires commercial products to be produced from commercial sources, and the commingling of public-service and commercial data is the firewall's primary failure mode.

Important Steps	Key Points	Reasons for Key Points
3. Verify the classification marking on the product.	Confirm that the Confidential marking is applied, and that the marking includes the customer identity and the product identifier.	Commercial products are restricted to the customer and the office, and the marking is the mechanism that enforces the restriction.
4. Verify the charge code on the production record.	Confirm that the commercial charge code was used, not the public-service charge code, and that the charge code matches the commercial arm's accounting.	The financial mechanism of the firewall requires separate charge codes, and the wrong charge code is a firewall breach.
5. Verify the customer identity and the agreed price against the commercial arrangements.	Confirm the customer identity against the office's commercial customer register, and confirm the price against the agreed pricing schedule.	The manager is the sole authority for commercial terms, and the verification prevents the release of products to unauthorized customers or at unauthorized prices.
6. Authorize the release by signing the commercial release log.	Record the product identifier, the customer identity, the price, the date, and the manager's signature in the commercial release log.	The commercial release log is the audit trail for the firewall, and the national audit authority examines the log during the annual audit.
7. Direct the production operator to release the product through the commercial channel.	Confirm that the release is through the commercial distribution channel, not the public-service distribution channel.	Commercial products must not enter the public-service distribution system, and the separation of channels is the distribution mechanism of the firewall.

13.5 The Monthly Reporting Cycle

The monthly reporting cycle is the procedure by which the office manager reports the office's performance to the desk chief. The report covers the office's production, its quality, its financial administration, and any incidents or exceptions. The report is submitted before the monthly review meeting, which the manager attends with the technical lead. The cycle is the office's primary upward reporting mechanism, and it is the mechanism by which the desk chief assesses the office's performance against the standards.

Important Steps	Key Points	Reasons for Key Points
1. Compile the monthly production statistics from the catalog and archive submodule.	Present the statistics by domain (flood, food, energy, minerals), not as an aggregate, and include the on-time delivery rate for each domain.	The desk chief needs domain-level visibility to assess whether the office is meeting its performance standards, and the aggregate rate hides domain-specific problems.
2. Compile the commercial arm financial summary.	Present the commercial revenue, the commercial costs, and the commercial surplus separately from the public-service expenditure, and present the revolving fund balance.	The firewall requires separate financial reporting, and the revolving fund balance is the indicator of the revenue-neutral model's sustainability.
3. Compile the quality metrics.	Include the classification compliance rate (which must be 100 percent), the base layer integrity check results, and the model validation results.	The classification compliance and the base layer integrity are non-negotiable standards, and any deviation is a performance failure that the desk chief must address.
4. Note any incidents or exceptions from the month.	Include both resolved and unresolved items, and for each incident, state the cause, the response, and the corrective action.	The desk chief needs to know what went wrong, not just what went right, and the unresolved items are the items that may require the desk chief's intervention.

Important Steps	Key Points	Reasons for Key Points
5. Submit the report to the desk chief before the monthly review meeting.	Submit at least two business days before the meeting, to give the desk chief time to read the report.	The desk chief cannot provide meaningful oversight if the report is received at the meeting, and the two-day lead time is the minimum for meaningful review.
6. Attend the monthly review meeting with the technical lead.	Bring the technical lead for technical questions, and do not answer technical questions without the technical lead's input.	The manager is accountable for the office's performance, but the technical lead is the officer who understands the technical details, and the manager who answers technical questions without the technical lead's input risks misinforming the desk chief.

13.6 The Incident Response Activation

The incident response activation is the procedure by which the office manager activates the office's response to a flood event or other emergency that requires geospatial intelligence support. The activation is triggered by an alert from the emergency management authority, by a sensor feed alert that indicates a developing event, or by a request from a consumer office. The procedure is the office's event-driven production mode, and it takes precedence over the scheduled cadence for the duration of the event.

Important Steps	Key Points	Reasons for Key Points
1. Receive the alert and acknowledge receipt to the emergency management authority.	Acknowledge within the timeframe specified in the inter-agency framework (typically 30 minutes), and state the office's initial response.	The emergency management authority needs to know that the office is responding, and the acknowledgement is the trigger for the authority's operational planning.
2. Assess the situation and determine the required products.	Identify which domains are affected, what products are needed, and what the delivery timeline is, in consultation with the hazard specialist.	The response must be proportionate to the event, and the products must be the products the consumer needs, not the products the office finds convenient to produce.

Important Steps	Key Points	Reasons for Key Points
3. Activate the hazard specialist and brief them on the event.	State the event, the required products, the delivery timeline, and the classification level, and confirm the specialist's understanding.	The specialist needs clear direction to produce the right intelligence, and the briefing is the mechanism that ensures the specialist and the manager share the same understanding of the task.
4. Notify the desk chief of the activation.	Report the activation, the initial assessment, and the planned response, and state whether the office needs additional resources.	The desk chief needs to know the office is in incident mode, and the notification is the mechanism that allows the desk chief to provide support or escalate if needed.
5. Monitor the production and verify product quality.	Conduct the base layer integrity check and the classification marking review even under time pressure; do not waive the checks for speed.	The base layer integrity check is non-negotiable, and intelligence produced against an unsynchronized base layer is wrong in ways that are not visible to the consumer.
6. Coordinate with the emergency management authority throughout the event.	Provide intelligence, not operational advice; the office's role is intelligence support, not operational command.	The emergency management authority is the operational authority, and the office that assumes operational command has exceeded its mandate and compromised its analytical independence.
7. Stand down the response and document the incident.	Record the timeline, the products produced, the issues encountered, and the lessons learned, and schedule the post-incident review.	The post-incident review depends on the documentation, and the documentation is the mechanism by which the office learns from the event and improves its future response.

13.7 The Manager's Discipline

The manager's discipline is the daily practice of maintaining the office's standards under operational pressure. The discipline has four components: production, firewall, classification, and

reporting. Each component is maintained through a specific procedure that surfaces the pressure and verifies the standard.

The production discipline covers the office's quality standards across the four domains: the base layer integrity check, the model validation protocols, the classification marking verification, and the archival integrity check. Three pressures bear on it. Production queue pressure arises when a product is overdue. Data quality pressure arises when the input data is incomplete or defective. Time pressure arises when a consumer needs the product urgently. The manager's production discipline is the refusal to release a product that has not passed the checks. The refusal is exercised at the manager's review gate: no product is released without the manager's review, and the review verifies that the checks were conducted.

The firewall discipline covers the separation between commercial and public-service work through the four mechanisms: separate charge codes, separate classification markings, the manager's sole authority for commercial release, and the commercial release log. Three pressures bear on it. Commercial pressure arises when a commercial customer needs the product urgently and the commercial data source is delayed. Workload pressure arises when the specialist is overloaded and the public-service data is available. Financial pressure arises when the commercial arm's surplus is below target. The manager's firewall discipline is the refusal to commingle the commercial and public-service data, charge codes, and classification markings. The refusal is exercised at the commercial release authorization (section 13.4), which verifies the firewall compliance before any commercial product is released.

The classification discipline covers the correct classification marking on every product. Three pressures bear on it. Handling burden pressure arises when a higher classification marking requires more handling, which slows the release. Consumer access pressure arises when the consumer needs the product but does not have the clearance for the higher marking. Default pressure arises when the lowest classification level is the easiest to apply. The manager's classification discipline is the application of the classification level that the content's sensitivity requires. The application is verified at the classification marking verification (section 16.4), which the production operator conducts and the manager reviews.

The reporting discipline covers the monthly report's timely submission, the inclusion of unresolved incidents, and the level of detail required for the desk chief's oversight. Three pressures bear on it. Time pressure arises when the report takes time to compile and the manager's day is full. Disclosure pressure arises when an unresolved incident is embarrassing to report. Simplification pressure arises when the desk chief appears not to need the detail. The manager's reporting discipline is the submission of the report on time, with the unresolved incidents included, and at the level of detail the desk chief needs. The submission is conducted through the monthly reporting cycle (section 13.5).

A concrete example illustrates the disciplines' interaction. A commercial customer requests a prospectivity product for a critical minerals commission, with a delivery deadline of two weeks. The resource economics specialist produces the product under the commercial charge code and presents the commercial release package to the manager. The manager reviews the package: the firewall compliance (the product was produced from commercial data sources), the classification marking (the Confidential marking with the customer identity), the charge code (the commercial charge code),

the customer identity (verified against the commercial customer register), and the agreed price (verified against the pricing schedule). The manager signs the commercial release log and directs the production operator to release the product through the commercial channel. The four disciplines are exercised in a single procedure: the production discipline (the product passed the validation), the firewall discipline (the commercial release authorization verified the compliance), the classification discipline (the Confidential marking was verified), and the reporting discipline (the commercial release log entry is the record for the monthly report).

The four disciplines are exercised through the manager's daily rhythm and the office's procedures. The consumers see the outcomes. The emergency management authority receives flood intelligence that is correctly marked and releasable. The strategic reserve authority receives the weekly crop condition bulletin on time. The commercial customers receive prospectivity products that are produced from commercial sources. The national audit authority sees the records: the commercial release log, the validation records, the classification marking verifications, and the monthly reports. The disciplines are established by the role family's training pipeline and qualifications framework, treated in Chapter 4, and maintained by the daily practice of the manager's work.

13.8 Worked Example: A Manager's Day During a Flood Event

The manager arrives at the office at 07:15 on a Wednesday morning. The sensor feed dashboard, monitored by the data pipeline operator, shows that the river stage at Gauge RB-04 on the Riverbourne has risen from 2.1 metres to 3.8 metres overnight, crossing the 5-year flood threshold of 3.5 metres. The rain gauge network shows 120 millimetres of rainfall in the preceding 6 hours. The data pipeline operator has already notified the technical lead, who has confirmed that the sensor feeds are active and the data is current.

At 07:30, the manager conducts an abbreviated morning production review (section 13.3). The review focuses on the flood event: the overnight sensor feed data, the satellite acquisitions (a Sentinel-1 SAR image was acquired at 05:30 UTC and is available for flood extent mapping), and the day's revised priorities. The manager cancels the scheduled production (the weekly crop condition bulletin, the monthly infrastructure monitoring cycle) and reprioritizes the office's work around the flood event response. The manager briefs the office: the top three priorities are (1) produce the flood inundation assessment for the Riverbourne Basin, (2) deliver the assessment to the emergency management authority within 4 hours, and (3) notify the desk chief of the activation.

At 07:45, the manager activates the incident response procedure (section 13.6). The manager acknowledges the event to the emergency management authority (within the 30-minute acknowledgement window), assesses the situation with the hazard specialist (the event requires a flood inundation map, an exposure report, and evacuation guidance), and briefs the specialist on the event, the required products, the delivery timeline (4 hours), and the classification level (Internal, because the exposure report includes critical infrastructure). The manager notifies the desk chief of the activation, reports the initial assessment, and states that the office does not need additional resources for this event.

At 08:00, the hazard specialist begins the flood model run (section 10.8). The data pipeline operator processes the Sentinel-1 SAR image and the LiDAR DEM. The technical lead validates the

processing output and the model run. The manager monitors the production dashboard, watching the model run's progress and the validation results, but does not interfere with the technical work. The manager's role during the model run is to be available: to answer the specialist's questions, to make decisions that the specialist cannot make (such as the classification level if the exposure report reveals unexpected critical infrastructure), and to coordinate with the emergency management authority.

At 10:30, the hazard specialist completes the model run and the validation. The production operator produces the flood inundation map (Internal classification), the exposure report (Internal), and the evacuation guidance (Public). The manager reviews the products at 11:00: the cartographic compliance (the blue depth gradient, the black parcel boundaries, the red infrastructure highlights), the classification marking (Internal for the map and the report, Public for the evacuation guidance), the format usability (GeoTIFF and PDF for the emergency management authority), and the provenance record (complete from the sensor feeds through the model run to the production). The manager authorizes the release at 11:15, within the 4-hour timeline.

At 11:30, the manager delivers the products to the emergency management authority through the secure data interface. The manager briefs the authority's duty officer on the assessment: the inundation extent, the 347 affected parcels, the estimated 1,200 people in the inundated area, and the 2 substations and 1 transmission corridor affected. The manager provides the intelligence, not operational advice: the manager does not recommend evacuation routes or shelter locations, because those decisions are the authority's responsibility.

At 14:00, the emergency management authority requests an updated assessment, based on the continued rainfall and the rising river stage. The manager directs the hazard specialist to produce an updated model run with the latest sensor feed data. The updated assessment is produced, reviewed, and delivered by 16:00. The manager coordinates with the authority throughout the afternoon, providing updates as the event evolves.

At 17:00, the river stage begins to recede. The manager stands down the incident response, documents the event (the timeline, the products produced, the issues encountered, the lessons learned), and schedules the post-incident review for the following week. The manager's day, which began at 07:15 with a sensor feed alert, ends at 17:30 with the documentation complete and the office returned to its scheduled cadence.

The day illustrates the manager's four disciplines operating together. The production discipline was maintained: the base layer integrity check and the classification marking review were conducted even under the 4-hour deadline. The firewall discipline was maintained: no commercial work was conducted during the event, and the commercial charge codes were not used. The classification discipline was maintained: the Internal marking was applied to the flood inundation map and the exposure report (because of the critical infrastructure exposure), and the Public marking was applied to the evacuation guidance. The reporting discipline was maintained: the desk chief was notified of the activation, and the event documentation will be included in the monthly report. The four disciplines, exercised through the manager's daily rhythm and the office's procedures, produced intelligence that was timely, correct, properly classified, and properly documented, and the intelligence served the emergency management authority's operational response. This single day also illustrates the manager's role as the office's single point of accountability: every decision that bore on the office's standards (the activation, the reprioritization, the classification level, the release

authorization, the stand-down) was made or confirmed by the manager, and every decision was documented for the audit trail. The manager who maintains this discipline under the pressure of a flood event is the manager whose office is trusted by the emergency management authority, by the desk chief, and by the national audit authority.

Chapter 14 — The Technical Lead

14.1 The Role of the Technical Lead

The technical lead is the geospatial office's senior technical professional. The lead is accountable for the office's pipeline configuration, the model validation, the systems integrity, and the specialist support. The position is the bridge between the operational tier and the office manager: the technical lead holds an operational rank but exercises a bridging function that includes technical configuration, validation, and the support of the specialists. The technical lead reports to the office manager, and performance is assessed by the manager on the office's technical performance.

The technical lead's accountability is described in operational detail through the job breakdown sheets in sections 14.3 through 14.6. The daily processing queue review (section 14.3) verifies that the overnight processing was conducted correctly. The model validation sampling (section 14.4) verifies that the specialists' model runs are producing accurate results. The weekly base layer integrity audit (section 14.5) verifies that the office's base layers are synchronized with the source authorities. The specialist support consultation (section 14.6) provides the technical support that the specialists require.

The technical lead's role is shaped by the pipeline architecture, treated in Part II. The pipeline is the shared backbone of the office's work, and the technical lead configures and maintains it. The technical lead configures the acquisition parameters in coordination with the data pipeline operator, the processing parameters, the modeling workbench, the production templates, and the archive configuration. The technical lead validates the processing outputs, the model runs, and the base layer integrity. The technical lead supports the specialists in their model runs, diagnosing and resolving the technical problems that the specialists encounter.

The technical lead's position is established at the operational tier of the role family, treated in Chapter 4, one rank above the operators and specialists. The position is held by an officer who has served as a specialist or as a data pipeline operator for a minimum period specified by the qualifications framework, holds the senior technical certification, and has completed the technical lead training pipeline. The position is a career position, and the appointment is made by the office manager with the desk chief's concurrence.

National mapping agencies organize the senior technical role in comparable ways. The United States Geological Survey's center chief scientists report to center directors and exercise authority for the center's scientific and technical work, including the data management and the quality assurance. Geoscience Australia's program science leads report to division heads and exercise authority for the program's scientific work, including the data management and the validation. IGN France's direction technique leads report to the director general and exercise authority for the agency's technical work, including the production standards and the quality assurance. The geospatial office's technical lead reports to the office manager and exercises authority for the office's technical work, including the pipeline configuration, the validation, and the specialist support. The pattern is consistent across the comparative practice: the senior technical officer reports to the office head and exercises authority for the technical work, including the validation and the quality assurance.

The technical lead's configuration management covers the standard parameter sets for the processing chain, the modeling workbench, the production templates, and the archive configuration.

The standard parameter sets are documented in the office's technical standards, and any deviation from the standard parameters requires the technical lead's authorization. The version control covers the model versions, the processing algorithm versions, and the base layer versions. The technical lead maintains the version register, which records the authorized versions and the dates of their authorization. The version control prevents the specialists from using untested or unauthorized versions, and the register is the record that the national audit authority examines during the annual audit. The version register records, for each model and each processing algorithm, the authorized version number, the authorization date, the reason for the authorization (a new release from the developer, a bug fix, a parameter recalibration), the testing results that justified the authorization, and the previous version that the new version replaces. When a new version is authorized, the technical lead conducts a regression test: the new version is run against a set of standard test cases (historical events with known outcomes), and the new version's results are compared with the previous version's results and with the known outcomes. If the new version's results are consistent with or better than the previous version's results, the new version is authorized. If the new version's results are worse, the new version is rejected, and the previous version remains authorized until the developer resolves the issue. The regression test is the quality gate that prevents the office from adopting a new version that produces worse results than the version it replaces, and the test is the operational practice that establishes the technical lead's configuration management discipline. The configuration management also covers the standard parameter sets for each domain. The hydrology plug-in's standard parameters include the Sacramento Soil Moisture Accounting model's parameters (the upper zone tension water capacity, the lower zone tension water capacity, the interflow recession constant, the baseflow recession constant), the HEC-RAS 2D model's parameters (the Manning's n roughness coefficients by land cover type, the time step, the mesh resolution), and the validation tolerance (0.30 metres depth RMSE, 80 per cent true positive rate for inundation extent). The crop model plug-in's standard parameters include the NDVI threshold values for each condition class (above-normal, normal, below-normal, poor, very poor), the yield model's regression coefficients (calibrated for each crop type and each region), and the validation tolerance (10 per cent yield accuracy at the district level). The change detection plug-in's standard parameters include the change threshold, the buffer widths by asset type, and the validation tolerance (85 per cent detection rate, 20 per cent false alarm rate). The prospectivity plug-in's standard parameters include the weights-of-evidence coefficients, the fuzzy logic membership functions, and the validation tolerance (70 per cent prospectivity accuracy, 3.0 prediction ratio). The standard parameter sets are documented in the office's technical standards, and any deviation requires the technical lead's authorization, with the deviation documented in the provenance record.

14.2 The Technical Lead's Daily Rhythm

The technical lead's day is organized around the office's fixed pipeline cadence. The day begins with the daily processing queue review (section 14.3), conducted after the morning production review and before the specialist support consultations. The review takes approximately thirty minutes and is conducted at the processing engine with the data pipeline operator present. The review verifies that the processing pipeline is operating correctly, that the processing parameters are appropriate for the

data being processed, and that the processing outputs meet the geometric and radiometric accuracy standards.

The model validation sampling (section 14.4) is conducted after the processing queue review. The sampling verifies that the specialists' model runs are producing accurate results by checking a sample of the model outputs against independent ground truth. The sampling is not a substitute for the specialists' own validation, which is conducted on every model run. The sampling is an independent verification that the specialists' validation is being conducted correctly and that the models are performing as expected. The sampling takes approximately one hour per day.

The specialist support consultations (section 14.6) are conducted throughout the day as needed. A specialist may request a consultation with the technical lead to diagnose a model problem, to configure a non-standard parameter set, or to resolve a data quality issue. The consultation is conducted at the modeling workbench or at the specialist's desk, and it is a working conversation rather than a formal meeting. The consultation is the mechanism by which the technical lead's technical knowledge is transferred to the specialists and by which the specialists' technical problems are resolved.

The weekly base layer integrity audit (section 14.5) is conducted on the first working day of each week. The audit verifies that the office's base layers are synchronized with the current state of the base layer sources. The audit is conducted in coordination with the data pipeline operator, who conducts the audit's operational steps, and the technical lead reviews the results and authorizes any updates. The audit takes approximately two hours on the day it is conducted.

The technical lead's rhythm is distinct from the manager's rhythm. The manager's rhythm is organized around the production, the firewall, and the reporting. The technical lead's rhythm is organized around the processing, the validation, and the support. The two rhythms are interdependent: the manager relies on the technical lead's validation to authorize the production, and the technical lead relies on the manager's direction to prioritize the work.

The daily rhythm is the technical lead's mechanism for maintaining technical situational awareness. A technical lead who skips the processing queue review does not know whether the overnight processing was conducted correctly. A technical lead who skips the model validation sampling does not know whether the specialists' models are performing as expected. A technical lead who skips the base layer integrity audit does not know whether the office's base layers are current. The discipline is the daily practice of being present at the processing engine and the modeling workbench and informed about the technical work.

The technical lead's weekly rhythm includes the base layer integrity audit, conducted on the first working day of each week with the data pipeline operator. The monthly rhythm includes the technical standards review, conducted on the first working day of each month, which verifies that the office's standard parameter sets are current with the source authorities' recommendations and the satellite operators' calibration updates. The monthly review also covers the model version register, which records the authorized model versions and the dates of their authorization. The technical lead's weekly and monthly rhythms complement the daily rhythm and ensure that the office's technical configuration is maintained over time.

14.3 The Daily Processing Queue Review

The daily processing queue review is the technical lead's first technical task of the day. It is conducted after the morning production review and before the specialist support consultations. The review verifies that the processing pipeline is operating correctly, that the processing parameters are appropriate for the data being processed, and that the processing outputs meet the geometric and radiometric accuracy standards. The review is conducted at the processing engine with the data pipeline operator present, and it takes approximately thirty minutes.

Important Steps	Key Points	Reasons for Key Points
1. Open the processing engine dashboard and review the overnight processing queue.	Check the queue status for each job (completed, failed, pending) and note any jobs that ran longer than the expected duration.	Jobs that run longer than expected may indicate a processing parameter problem or a data quality issue, and the issue must be investigated before the processed data enters the modeling workbench.
2. Verify the processing parameters for each completed job.	Confirm that the georeferencing, calibration, mosaicking, and cloud masking parameters match the standard parameter set for the data type and the intended domain.	Non-standard parameters produce non-standard products, and the non-standard products may not be compatible with the modeling workbench's expectations.
3. Review the geometric accuracy validation results for each completed job.	Check the root mean square error against the tolerance, and flag any jobs where the error exceeds the tolerance.	The geometric accuracy is the integrity guarantee of the processing function, and products that fail the geometric validation must not enter the modeling workbench.
4. Review the radiometric calibration results for each completed job.	Verify that the calibration coefficients were applied correctly, and check the radiometric consistency across overlapping scenes.	Radiometric inconsistency propagates through the modeling workbench and produces inconsistent analytical results, particularly in the food security domain's vegetation indices.

Important Steps	Key Points	Reasons for Key Points
5. Investigate any failed jobs and determine the cause.	Distinguish between data quality failures (the input data was defective) and processing parameter failures (the parameters were wrong for the data), and correct the appropriate cause.	Re-running a failed job without correcting the cause reproduces the failure, and the correction requires knowing whether the problem is in the data or the parameters.
6. Release the validated processing outputs to the modeling workbench.	Confirm that each released output has a complete provenance record, including the source data, the processing parameters, and the validation results.	The provenance chain is the integrity guarantee of the pipeline, and a processing output without a complete provenance record is a break in the chain.

14.4 The Model Validation Sampling

The model validation sampling is the technical lead's daily quality assurance task, conducted after the processing queue review. The sampling verifies that the specialists' model runs are producing accurate results by checking a sample of the model outputs against independent ground truth. The specialists conduct their own validation on every model run. The sampling is an independent verification that the specialists' validation is being conducted correctly and that the models are performing as expected.

Important Steps	Key Points	Reasons for Key Points
1. Select a sample of model runs from the past 24 hours for validation review.	Select at least one run from each domain, and include any runs that the specialist flagged as marginal or that used non-standard parameters.	The sample must be representative of the office's work, and the marginal and non-standard runs are the runs most likely to have validation issues.
2. Retrieve the model run's provenance record from the catalog and archive submodule.	Verify that the provenance record is complete, including the source data, the processing parameters, the model version, the model parameters, and the validation results.	An incomplete provenance record is a break in the chain, and the model run cannot be validated without the complete record.

Important Steps	Key Points	Reasons for Key Points
3. Review the specialist's validation results.	Check the accuracy metrics against the domain's tolerance, and verify that the validation was conducted against independent ground truth, not against the data used to calibrate the model.	Validation against the calibration data is not validation; it is circular reasoning, and it produces accuracy metrics that overstate the model's performance.
4. Conduct an independent validation on the sampled runs.	Use a different set of ground truth from the specialist's validation set, and calculate the accuracy metrics using the standard formulas.	The independent validation is the check on the specialist's validation, and the use of a different ground truth set ensures that the check is genuinely independent.
5. Compare the independent validation results with the specialist's results.	If the results are consistent, the specialist's validation is confirmed; if the results differ, investigate the cause of the discrepancy.	Consistent results confirm that the validation is being conducted correctly; differing results indicate either a validation error or a model problem, and the investigation determines which.
6. Record the sampling results in the catalog and archive submodule.	Record the sampled runs, the independent validation results, the comparison with the specialist's results, and any discrepancies and their resolutions.	The sampling record is the quality assurance record of the modeling function, and it is the record that the national audit authority examines during the annual audit.

14.5 The Base Layer Integrity Audit

The base layer integrity audit is the technical lead's weekly task, conducted on the first working day of each week. The audit verifies that the office's base layers are synchronized with the current state of the base layer sources: the cadastre feed from the cadastral authority, the elevation data from the national mapping authority, the land cover data from the national land cover authority, the hydrography from the national hydrological service, and the administrative boundaries from the national mapping authority. The audit is conducted in coordination with the data pipeline operator, who conducts the operational steps, and the technical lead reviews the results and authorizes any updates.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the current version of each base layer from the spatial data infrastructure submodule.	Record the version number, the version date, and the source for each base layer.	The version record is the baseline against which the synchronization is checked, and the version must be recorded before the check is conducted.
2. Retrieve the current version of each base layer from the source authority.	Query the source authority's data interface for the current version, and record the version number and the version date.	The source authority's version is the authoritative version, and the office's version must match it.
3. Compare the office's version with the source authority's version for each base layer.	If the versions match, the base layer is synchronized; if the versions do not match, the base layer is outdated and must be updated before any production run that depends on it.	Outdated base layers produce intelligence that is wrong in ways that are not visible to the consumer, and the comparison is the check that prevents the wrongness.
4. Update any outdated base layers.	Download the current version from the source authority, ingest it into the spatial data infrastructure, and verify the ingestion's integrity.	The update is the correction of the outdated base layer, and the verification ensures that the update was conducted correctly.
5. Record the audit results in the catalog and archive submodule.	Record the base layers checked, the versions compared, the discrepancies found, and the updates conducted.	The audit record is the compliance record for the base layer integrity principle, and it is the record that the national audit authority examines during the annual audit.

14.6 The Specialist Support Consultation

The specialist support consultation is the technical lead's daily task of providing technical support to the specialists. The consultation is conducted throughout the day as needed, and it covers any technical issue that the specialists encounter in their model runs. The consultation is a working conversation between the technical lead and the specialist, conducted at the modeling workbench or at the specialist's desk. The consultation is the mechanism by which the technical lead's technical knowledge is transferred to the specialists and by which the specialists' technical problems are resolved.

Important Steps	Key Points	Reasons for Key Points
1. Receive the specialist's request for support.	Ask the specialist to state the problem clearly, including the model, the data, the parameters, and the unexpected result.	A clear problem statement is the foundation of the diagnosis, and the specialist's initial statement often contains the information needed to identify the cause.
2. Review the model run's provenance record with the specialist.	Trace the provenance chain from the source data through the processing, the modeling, and the production, and identify any breaks or anomalies in the chain.	The provenance chain is the integrity guarantee of the pipeline, and breaks or anomalies in the chain are the most common cause of unexpected model results.
3. Diagnose the problem in consultation with the specialist.	Distinguish between data problems (the input data was defective or inappropriate), parameter problems (the parameters were wrong for the data or the model), and model problems (the model itself has a defect), and correct the appropriate cause.	The correction depends on the diagnosis, and the wrong diagnosis produces the wrong correction, which may not fix the problem and may introduce new problems.
4. Assist the specialist in implementing the correction.	Do not implement the correction for the specialist; guide the specialist through the implementation, so that the specialist learns the correction and can apply it independently in the future.	The consultation is a training opportunity as well as a problem-solving opportunity, and the specialist who learns the correction becomes a more capable officer.

Important Steps	Key Points	Reasons for Key Points
5. Document the problem, the diagnosis, and the correction in the office's technical log.	Record the problem statement, the diagnosis, the correction, and the specialist's name, and file the record in the office's technical log.	The technical log is the office's institutional memory for technical problems and their solutions, and it is the resource that the office consults when the same problem arises again.

14.7 The Technical Lead's Discipline

The technical lead's discipline is the daily practice of maintaining the office's technical integrity under operational pressure. The discipline has four components: processing, validation, base layer, and support. Each component is maintained through a specific procedure that surfaces the pressure and verifies the standard.

The processing discipline covers the processing standards: the georeferencing, calibration, mosaicking, and cloud masking parameter sets; the geometric accuracy tolerance; and the radiometric consistency requirements. Three pressures bear on it. Throughput pressure arises when the queue is full and the parameters are time-consuming to verify. Data quality pressure arises when the input data is defective and the standard parameters produce poor results. Consumer urgency pressure arises when the specialist needs the data urgently. The technical lead's processing discipline is the refusal to release a processing output that has not passed the validation. The refusal is exercised at the daily processing queue review (section 14.3), which is the gate through which the processed data enters the modeling workbench.

The validation discipline covers the validation protocols: the model validation against independent ground truth, the accuracy metrics calculation, the tolerance comparison, and the validation record. Three pressures bear on it. Time pressure arises when validation takes time and the consumer needs the product. Specialist confidence pressure arises when the specialist is confident in the model and resents the independent check. Data availability pressure arises when the ground truth is not available for the current run. The technical lead's validation discipline is the conduct of the validation on every run. The conduct is exercised at the model validation sampling (section 14.4), which is the independent check on the specialists' validation.

The base layer discipline covers the synchronization of the office's base layers with the source authorities' current versions. Three pressures bear on it. Schedule pressure arises when the weekly audit takes time and the week is busy. Discrepancy pressure arises when a base layer is outdated and the update requires re-running the affected models. Source authority pressure arises when the source authority's interface is unreliable. The technical lead's base layer discipline is the conduct of the weekly audit. The conduct is exercised at the base layer integrity audit (section 14.5), which is the gate through which the office's base layers are verified.

The support discipline covers the technical support that the specialists require. Three pressures bear on it. Time pressure arises when the technical lead's day is full. Diagnostic pressure arises when the specialist's problem is difficult to diagnose. Training pressure arises when the specialist needs to

learn the correction, not have it done for them. The technical lead's support discipline is the provision of the support that the specialists require. The provision is exercised at the specialist support consultation (section 14.6), which is the working conversation through which the technical lead diagnoses and resolves the specialist's problems.

A concrete example illustrates the disciplines' interaction. The hazard specialist requests a consultation: the flood model run for a monitored basin produced an inundation extent larger than the previous run's extent, and the specialist cannot identify the cause. The technical lead reviews the model run's provenance record with the specialist, tracing the chain from the source data through the processing and the modeling. The technical lead identifies a radiometric calibration inconsistency in the processed imagery: the calibration coefficients for one scene were applied incorrectly, the resulting radiometric inconsistency propagated through the vegetation index calculation, and the affected land cover classification changed the hydraulic model's roughness parameters. The technical lead guides the specialist through the correction: the scene is reprocessed with the correct coefficients, the vegetation index is recalculated, the land cover classification is updated, and the flood model is re-run. The four disciplines are exercised in a single consultation: the processing discipline (the radiometric calibration was verified), the validation discipline (the re-run was validated against the observed event), the base layer discipline (the land cover classification was verified), and the support discipline (the specialist learned the correction).

The four disciplines are exercised through the technical lead's daily rhythm and the office's procedures. The specialists receive the validated processing outputs and the current base layers they need for their model runs. The office manager authorizes the production on the basis of the technical lead's validation. The national audit authority examines the validation records and the base layer integrity records during the annual audit. The disciplines are established by the role family's training pipeline and qualifications framework, treated in Chapter 4, and maintained by the daily practice of the technical lead's work.

14.8 Worked Example: Diagnosing a Model Validation Failure

The hazard specialist submits a flood model run for the Riverbourne Basin to the technical lead for review. The model run used the standard parameters: HEC-RAS 2D version 6.3.1, the LiDAR DEM at 2-metre resolution (resampled to 5 metres for the mesh), the Manning's n values from the land cover (0.035 for the main channel, 0.060 for urban, 0.080 for forest, 0.030 for agriculture), and the event-specific rainfall from the 15 rain gauges. The specialist's validation shows a peak stage at Gauge RB-04 of 3.95 metres against the observed 3.85 metres (a difference of 0.10 metres, within the 0.30-metre tolerance), but the inundation extent true positive rate is 72 per cent, below the 80 per cent tolerance. The specialist has flagged the run as marginal and has requested the technical lead's review.

The technical lead retrieves the model run's provenance record from the catalog and archive submodule. The provenance record shows the complete chain: the Sentinel-1 SAR image (acquisition ID S1B_IW_SLC__1SDV_20240117T053000), the LiDAR DEM (version 2019-1, acquired 2019, 2-metre resolution), the land cover (National Land Cover Database 2021, 30-metre resolution), the hydrography (National Hydrography Dataset, version 2023-10), the rain gauge data (15 gauges, 15-minute interval), and the river stage data (8 gauges, 15-minute interval). The chain is complete; there are no breaks.

The technical lead reviews the specialist's validation methodology. The specialist validated the model against the observed river stages at the 8 gauges and against the Sentinel-1 SAR-derived flood extent. The gauge validation is within tolerance (the peak stages at all 8 gauges are within 0.30 metres of the observed values). The extent validation is below tolerance: the model's predicted inundation extent covers a larger area than the SAR-observed extent, with the over-prediction concentrated in the southeastern portion of the basin, near the confluence of the Riverbourne and the Millbrook tributary.

The technical lead investigates the over-prediction. The technical lead examines the DEM in the southeastern portion of the basin and discovers that the LiDAR DEM, acquired in 2019, does not reflect a flood defence embankment that was constructed in 2022 along the Millbrook tributary. The embankment is approximately 1.5 metres high and 800 metres long, and it was constructed by the water authority as part of a flood risk reduction project. The DEM shows the pre-embankment terrain, and the model's inundation extends over the embankment's location because the model does not see the embankment.

The technical lead diagnoses the problem as a data problem: the DEM is outdated. The embankment was constructed after the DEM was acquired, and the DEM does not reflect the current ground surface. The model's inundation extent is too large because the model allows the flood water to flow over the area where the embankment now prevents the flow. The specialist's model run is analytically correct (the model was run with the correct parameters and the correct methods), but the input data (the DEM) is outdated, and the outdated data produces an inaccurate result.

The technical lead does not correct the DEM (the DEM is maintained by the national mapping authority, and the correction requires a new LiDAR acquisition). Instead, the technical lead directs the specialist to document the discrepancy in the model run's provenance record: the DEM does not reflect the 2022 Millbrook embankment, and the model's inundation extent is over-predicted in the embankment's vicinity. The technical lead reports the discrepancy to the national mapping authority through the formal discrepancy reporting channel (treated in Chapter 8, section 8.8), requesting a new LiDAR acquisition for the southeastern portion of the basin. The technical lead also directs the specialist to note the over-prediction in the flood inundation map's metadata, so that the emergency management authority is aware of the limitation when using the map for evacuation planning.

The technical lead then conducts an independent validation of the model run, excluding the southeastern portion where the over-prediction is known. The independent validation uses a different set of check points from the specialist's validation set, and the results confirm the specialist's gauge validation (all check points within 0.30 metres) and show an inundation extent true positive rate of 84 per cent in the non-affected areas (above the 80 per cent tolerance). The technical lead records the sampling results in the catalog and archive submodule, with the note that the overall true positive rate (72 per cent) is below tolerance due to the known DEM discrepancy in the southeastern portion, and that the true positive rate in the non-affected areas (84 per cent) is within tolerance.

The technical lead authorizes the release of the flood inundation map with the metadata note about the southeastern over-prediction, and the office manager reviews the product and authorizes the release to the emergency management authority. The release is documented in the provenance record, with the technical lead's validation note and the metadata caveat. The discrepancy report to the national mapping authority is the mechanism by which the DEM will eventually be updated, and the metadata caveat is the mechanism by which the consumer is informed of the limitation in the interim.

This worked example illustrates the technical lead's four disciplines operating together. The processing discipline was exercised (the processing parameters were verified). The validation discipline was exercised (the independent validation confirmed the specialist's validation and identified the cause of the over-prediction). The base layer discipline was exercised (the DEM discrepancy was identified and reported). The support discipline was exercised (the specialist was guided through the diagnosis and the documentation). The four disciplines, exercised through the technical lead's daily rhythm and the office's procedures, produced intelligence that was technically sound, honestly documented, and properly caveated, and the intelligence served the emergency management authority's operational response while the DEM discrepancy was being resolved. The example also illustrates the technical lead's role as the office's technical problem-solver: the specialist identified the symptom (the low true positive rate), but the technical lead identified the cause (the outdated DEM), determined the response (document the limitation, report the discrepancy, release with caveat), and conducted the independent validation that confirmed the model's performance in the non-affected areas. The technical lead who maintains this diagnostic discipline is the technical lead whose office's models are trusted by the consumers and whose technical standards are verified by the national audit authority.

Chapter 15 — The Data Pipeline Operator

15.1 The Role of the Data Pipeline Operator

The data pipeline operator is the geospatial office's data intake and processing operator. The operator is accountable for the daily ingestion, the processing queue management, the base layer maintenance, and the license tracking. The operator is the primary operator of the acquisition and processing submodules of the geospatial information system, treated in Chapter 29. The operator reports to the technical lead, who configures the pipeline parameters and validates the processing outputs, and performance is assessed by the technical lead on the pipeline's operational performance.

The operator's work is the front end of the pipeline. Every product that the office delivers passes through the operator's hands at the acquisition and processing stages. The quality of the operator's work determines the quality of every subsequent product. If the operator ingests the wrong data, or ingests the right data with the wrong parameters, or fails to record the license terms, the defect propagates through the pipeline and emerges in the product. The operator's work is invisible to the consumer, but it is the foundation of the office's intelligence.

The operator's role is shaped by the pipeline architecture, treated in Part II. The pipeline is the shared backbone of the office's work, and the operator operates it on a daily basis. The operator operates the acquisition desk, ingesting the data that the specialists request and the office subscribes to. The operator operates the processing engine, running the processing chain that transforms raw data into analysis-ready products. The operator operates the spatial data infrastructure, maintaining the base layers that underlie every analysis. The operator operates the license management function, recording and enforcing the data license terms.

The operator's position is established at the operational tier of the role family, treated in Chapter 4. The position is held by an officer who has completed the training pipeline, holds the operator certification, and has served a probationary period under the supervision of a senior operator. The position is a career position, and the operator may progress to the senior operator rank and then to the technical lead position with the additional certifications and experience specified by the qualifications framework. The operator's training pipeline includes the foundational course (covering the pipeline architecture, the information system, the classification framework, and the office's procedures), the acquisition desk course (covering the satellite tasking, the imagery purchasing, the sensor feed ingestion, and the license management), the processing engine course (covering the georeferencing, the calibration, the mosaicking, the projection management, the cloud masking, and the data fusion), and the spatial data infrastructure course (covering the base layer maintenance, the synchronization check, and the discrepancy reporting). The operator must pass the certification assessment for each course before operating the corresponding submodule independently. The certification is maintained through continuing education, with refresher courses and update courses required on a periodic basis to keep the operator's knowledge current with the evolving tools and methods.

National mapping agencies organize the data operations function in comparable ways. The United States Geological Survey's Earth Resources Observation and Science Center employs data operations specialists who operate the Landsat and Sentinel ingestion and processing pipelines under the supervision of the center's data management team. The European Space Agency's Copernicus

Data and Information Access Services employs operations specialists who operate the Sentinel ingestion and processing pipelines under the supervision of the mission management team. Geoscience Australia's National Earth and Marine Observations branch employs operations specialists who operate the satellite data ingestion and processing pipelines under the supervision of the branch's data management team. The geospatial office's data pipeline operator operates the ingestion, processing, and base layer maintenance pipelines under the supervision of the technical lead, with the operational steps conducted to the technical lead's configuration and the outputs validated by the technical lead before release to the modeling workbench.

The operator's operational responsibilities are specified in the job breakdown sheets in sections 15.3 through 15.6. The morning ingestion check (section 15.3) verifies that all overnight data acquisitions were received and ingested correctly. The processing pipeline configuration (section 15.4) sets up the day's processing runs. The weekly base layer integrity audit (section 15.5) verifies the base layer synchronization. The license compliance review (section 15.6) verifies the office's adherence to the data license terms. The operator conducts the operational steps of each procedure, and the technical lead reviews the results and authorizes any deviations or corrections.

15.2 The Operator's Daily Rhythm

The data pipeline operator's day is organized around the office's fixed pipeline cadence. The day begins with the morning ingestion check (section 15.3), conducted before any other activity and before the morning production review. The check verifies that all overnight data acquisitions were received and ingested correctly. The check is conducted at the acquisition desk and takes approximately thirty minutes. The check is the operational foundation of the day's pipeline.

The processing pipeline configuration (section 15.4) follows the morning ingestion check. The configuration selects the processing parameters, submits the processing jobs to the processing engine, and monitors the jobs' execution. The configuration is conducted under the technical lead's oversight, and the operator who configures the pipeline correctly produces analysis-ready data that the specialists can use. The configuration takes approximately two hours per day, depending on the volume of data to be processed.

The license compliance review (section 15.6) is conducted on a rolling basis as new data is acquired, as products are produced, and as the office's license terms are updated. The review verifies that the office's use of its acquired data complies with the data license terms. The review is conducted at the catalog and archive submodule, where the license records are maintained. The review takes approximately one hour per day on a rolling basis.

The weekly base layer integrity audit (section 15.5) is conducted on the first working day of each week, in coordination with the technical lead. The audit verifies that the office's base layers are synchronized with the source authorities' current versions. The operator conducts the audit's operational steps, and the technical lead reviews the results and authorizes any updates. The audit takes approximately two hours on the day it is conducted.

The operator's rhythm is distinct from the specialists' rhythm. The specialists operate on variable cadences determined by their domains: the hazard specialist on the flood and energy schedules, the resource economics specialist on the food and minerals schedules. The operator operates on the fixed cadence, which serves all four domains through the shared pipeline. If the operator's rhythm falters —

a missed ingestion, a delayed processing run, an outdated base layer — all four domains falter. The operator's rhythm is the heartbeat of the office, and the specialists' variable cadences depend on it.

The daily rhythm is the operator's mechanism for maintaining the pipeline's operational reliability. An operator who skips the morning ingestion check does not know whether the overnight acquisitions were received. An operator who configures the processing pipeline incorrectly does not produce analysis-ready data. An operator who skips the license compliance review does not know whether the office is in compliance with the data license terms. The discipline is the daily practice of being present at the acquisition desk and the processing engine and informed about the pipeline's state.

The operator's weekly rhythm includes the base layer integrity audit, conducted on the first working day of each week in coordination with the technical lead. The monthly rhythm includes the archive audit, conducted on the first working day of each month in coordination with the production operator, and the license subscription review, which verifies that the office's data subscriptions are current and that the renewals are scheduled. The operator's monthly rhythm includes the data subscription review, conducted on the first working day of each month, which verifies that the office's data subscriptions (the Copernicus Open Access Hub subscription, the Landsat subscription, the commercial imagery framework agreements, and the sensor feed subscriptions) are active, that the subscription terms have not changed, and that the renewal dates are scheduled. The subscription review is the mechanism by which the operator prevents the disruption that would occur if a subscription expired without renewal. The monthly rhythm also includes the processing engine maintenance check, conducted in coordination with the technical lead, which verifies that the processing engine's software is up to date, that the hardware is performing within specifications, and that the storage capacity is sufficient for the projected acquisition volume. The maintenance check is the mechanism by which the operator prevents the processing failures that would occur if the engine's software was outdated or the storage was full. The operator's weekly and monthly rhythms complement the daily rhythm and ensure that the pipeline's operational reliability is maintained over time.

15.3 The Morning Ingestion Check

The morning ingestion check is the data pipeline operator's first task of the day. It is conducted before any other activity and before the morning production review. The check verifies that all overnight data acquisitions were received and ingested correctly, and it identifies any acquisitions that failed or arrived late. The check is conducted at the acquisition desk, and it takes approximately thirty minutes. The check is the operational foundation of the day's pipeline.

Important Steps	Key Points	Reasons for Key Points
1. Open the acquisition desk dashboard and review the overnight acquisition log.	Check the log for each scheduled acquisition: satellite tasking, imagery purchases, and sensor feed subscriptions.	The acquisition log is the record of what was expected, and the operator needs to know what was expected before checking what was received.
2. Verify that each scheduled satellite acquisition was received.	Cross-check the acquisition schedule against the ingestion log, and flag any missing acquisitions immediately to the technical lead.	Missing acquisitions break the pipeline for the affected domain, and the technical lead needs to know before the specialist starts a model run on incomplete data.
3. Verify that each sensor feed subscription is active and delivering data.	Check the latest data timestamp for each sensor feed; a feed that has not updated in the expected interval may be down.	Sensor feeds are the real-time data source for the flood risk domain, and a feed that is down silently produces intelligence that is based on stale data.
4. Verify that each received acquisition was ingested correctly.	Check the ingestion log for each acquisition: the file integrity check, the format verification, and the metadata extraction.	An acquisition that was received but not ingested correctly is as useless as an acquisition that was not received, and the ingestion log is the record that distinguishes the two.
5. Record the license terms for each new acquisition in the catalog and archive submodule.	Record the data source, the license type, the license restrictions, and the license expiration date for each acquisition.	The license record is the legal integrity guarantee of the pipeline, and the record must be made at the moment of acquisition, not later, to prevent omissions.

Important Steps	Key Points	Reasons for Key Points
6. Notify the technical lead of any missing acquisitions, failed ingestions, or sensor feed outages.	Provide a clear summary of the issue, the affected domain, and the recommended corrective action.	The technical lead is the officer who can authorize corrective action, and the operator's summary is the information the technical lead needs to act.
7. Release the ingested data to the processing queue.	Confirm that each released dataset has a complete provenance record, including the source, the license, and the ingestion metadata.	The provenance chain begins at acquisition, and a dataset without a complete provenance record is a break in the chain that propagates through the entire pipeline.

15.4 The Processing Pipeline Configuration

The processing pipeline configuration is the data pipeline operator's daily task of setting up the processing runs for the day's ingested data. The configuration selects the processing parameters, submits the processing jobs to the processing engine, and monitors the jobs' execution. The configuration is conducted under the technical lead's oversight. The operator who configures the pipeline correctly produces analysis-ready data that the specialists can use; the operator who configures it incorrectly produces data that fails validation or produces inaccurate models.

Important Steps	Key Points	Reasons for Key Points
1. Review the ingested data in the processing queue and determine the required processing chain for each dataset.	Identify the data type (optical, radar, elevation, vector), the intended domain consumer, and the standard processing parameter set for the combination.	The processing chain is determined by the data type and the intended use, and the wrong chain produces data that is inappropriate for the specialist's model.
2. Configure the georeferencing parameters for each dataset.	Select the ground control point library appropriate for the data's geographic coverage, and set the transformation model (polynomial order or spline) per the standard parameter set.	The georeferencing parameters determine the geometric accuracy of the processed product, and the wrong parameters produce geometrically distorted products.

Important Steps	Key Points	Reasons for Key Points
3. Configure the calibration parameters for each dataset.	Select the calibration coefficients provided by the satellite operator, and set the calibration target (top-of-atmosphere reflectance or surface reflectance) per the standard parameter set.	The calibration parameters determine the radiometric accuracy, and the wrong parameters produce radiometrically inconsistent products that corrupt the vegetation indices.
4. Configure the mosaicking parameters for datasets that require mosaicking.	Set the color balancing method, the seam feathering width, and the output mosaic extent per the standard parameter set.	The mosaicking parameters determine the spatial continuity of the processed product, and the wrong parameters produce mosaics with visible seams or color inconsistencies.
5. Configure the cloud masking parameters for optical datasets.	Select the cloud masking algorithm (the processing authority's standard algorithm or the office's supplemental algorithm) and set the confidence threshold.	The cloud masking parameters determine the cloud contamination of the processed product, and the wrong parameters produce products that are contaminated by cloud artifacts.
6. Submit the processing jobs to the processing engine.	Submit the jobs in the priority order determined by the technical lead, and verify that each job's submission was acknowledged by the processing engine.	The priority order ensures that the most urgent processing is conducted first, and the acknowledgement verifies that the job entered the processing queue.
7. Monitor the processing jobs' execution and report any failures.	Check the job status at regular intervals, and when a job fails, note the failure reason and report it to the technical lead before re-submitting.	Re-submitting a failed job without understanding the failure reproduces the failure, and the technical lead is the officer who can diagnose and correct the cause.

15.5 The Weekly Base Layer Integrity Audit

The weekly base layer integrity audit is the data pipeline operator's weekly task, conducted on the first working day of each week in coordination with the technical lead. The audit verifies that the office's base layers are synchronized with the current state of the base layer sources. The operator conducts the audit's operational steps, and the technical lead reviews the results and authorizes any updates. The audit is conducted at the spatial data infrastructure submodule.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the current version of each base layer from the spatial data infrastructure submodule.	Record the version number, the version date, and the source authority for each base layer: elevation, land cover, hydrography, administrative boundaries, and cadastre.	The version record is the baseline against which the synchronization is checked, and the record must be complete before the comparison is conducted.
2. Query each source authority's data interface for the current version.	Query the national mapping authority for elevation and administrative boundaries, the national land cover authority for land cover, the national hydrological service for hydrography, and the cadastral authority for the cadastre feed.	The source authority's version is the authoritative version, and the office's version must match it for the base layer to be considered synchronized.
3. Compare the office's version with the source authority's version for each base layer.	If the versions match, the base layer is synchronized; if the versions do not match, the base layer is outdated and must be updated.	Outdated base layers produce intelligence that is wrong in ways that are not visible to the consumer, and the comparison is the check that prevents the wrongness.
4. Download the current version of any outdated base layer from the source authority.	Download the full updated dataset, not a delta or a patch, unless the source authority provides a reliable delta mechanism.	A full download ensures that the office has the complete current version, and deltas or patches may introduce inconsistencies if they are not applied correctly.

Important Steps	Key Points	Reasons for Key Points
5. Ingest the updated base layer into the spatial data infrastructure submodule.	Run the ingestion with the integrity verification enabled, and verify that the ingestion's checksum matches the source authority's checksum.	The integrity verification ensures that the download was not corrupted, and the checksum comparison is the verification that the ingested data matches the source.
6. Report the audit results to the technical lead.	Provide the version comparison results, the updates conducted, and any discrepancies that could not be resolved.	The technical lead is the officer who authorizes the updates and who needs to know about any unresolvable discrepancies, and the report is the information the technical lead needs to act.
7. Record the audit results in the catalog and archive submodule.	Record the base layers checked, the versions compared, the discrepancies found, the updates conducted, and the date of the audit.	The audit record is the compliance record for the base layer integrity principle, and it is the record that the national audit authority examines during the annual audit.

15.6 The License Compliance Review

The license compliance review is the data pipeline operator's rolling task of verifying that the office's use of its acquired data complies with the data license terms. The review is conducted as new data is acquired, as products are produced, and as the office's license terms are updated. The review is conducted at the catalog and archive submodule, where the license records are maintained. The review prevents the office from releasing products that infringe the data providers' intellectual property.

Important Steps	Key Points	Reasons for Key Points
1. Review the license terms for each newly acquired dataset.	Identify the license type (free and open, commercial, restricted), the permitted uses, the prohibited uses, and the attribution requirements.	The license terms determine what the office may do with the data, and the operator needs to know the terms before the data enters the processing pipeline.

Important Steps	Key Points	Reasons for Key Points
2. Verify that the intended use of the dataset is permitted by the license terms.	Cross-check the intended domain consumer and the intended product type against the license's permitted uses, and flag any intended use that is not permitted.	Using data for a purpose not permitted by the license is a license breach, and the breach exposes the office to legal liability and reputational damage.
3. Verify that the attribution requirements are met for each product that uses the dataset.	Check the product's attribution block to confirm that the data source is credited in the form required by the license.	Missing or incorrect attribution is a license breach, and the breach is visible to the consumer and to the data provider.
4. Check the license expiration dates for the office's data subscriptions.	Flag any subscription that will expire within the next 30 days, and notify the office manager to arrange renewal.	An expired subscription means the office loses access to the data, and the loss breaks the pipeline for the affected domain.
5. Review the commercial imagery licenses for the commercial minerals products.	Verify that the commercial minerals products do not use commercial imagery licensed for public-service use only, and that the commercial charge code matches the commercial license.	The firewall requires commercial products to be produced from commercial sources, and the use of public-service-licensed imagery for commercial products is a firewall breach as well as a license breach.
6. Record the license compliance review results in the catalog and archive submodule.	Record the datasets reviewed, the license terms verified, the attributions checked, and any compliance issues identified and resolved.	The compliance record is the legal integrity guarantee of the pipeline, and it is the record that the national audit authority examines during the annual audit.

15.7 The Operator's Discipline

The data pipeline operator's discipline is the daily practice of maintaining the pipeline's operational reliability under operational pressure. The discipline has four components: ingestion, processing, base layer, and license. Each component is maintained through a specific procedure that surfaces the pressure and verifies the standard.

The ingestion discipline covers the ingestion completeness: every scheduled acquisition received and ingested correctly, with the file integrity check, the format verification, the metadata

extraction, and the license term recording conducted at the moment of acquisition. Three pressures bear on it. Volume pressure arises when the overnight acquisition volume is large. Source authority pressure arises when the source authority's interface is unreliable. Time pressure arises when the morning ingestion check takes time and the day's processing is waiting. The operator's ingestion discipline is the conduct of the check thoroughly, regardless of the volume or the interface. The conduct is exercised at the morning ingestion check (section 15.3), which is the gate through which the day's data enters the pipeline.

The processing discipline covers the processing standards: the georeferencing, calibration, mosaicking, and cloud masking parameter sets, which are configured by the technical lead and applied by the operator. Three pressures bear on it. Parameter pressure arises when the standard parameters are wrong for the data and the operator is tempted to deviate. Throughput pressure arises when the queue is full and the operator is tempted to skip the parameter verification. Consumer urgency pressure arises when the specialist needs the data urgently. The operator's processing discipline is the application of the standard parameters and the reporting of any deviation to the technical lead. The application is exercised at the processing pipeline configuration (section 15.4), which is the gate through which the ingested data enters the processing engine.

The base layer discipline covers the synchronization of the office's base layers with the source authorities' current versions. Three pressures bear on it. Schedule pressure arises when the weekly audit takes time and the week is busy. Discrepancy pressure arises when a base layer is outdated and the update requires re-running the affected models. Source authority pressure arises when the source authority's interface is unreliable. The operator's base layer discipline is the conduct of the weekly audit. The conduct is exercised at the base layer integrity audit (section 15.5), which is the gate through which the office's base layers are verified.

The license discipline covers the office's adherence to the data license terms that govern the use of the data the office acquires. Three pressures bear on it. Volume pressure arises when the license records are many and the review takes time. Attribution pressure arises when the attribution block is easy to omit and the consumer does not check it. Firewall pressure arises when the commercial minerals product is waiting and the public-service-licensed imagery is available. The operator's license discipline is the recording of the license terms at the moment of acquisition, the verification of the intended use against the license, and the reporting of any breach to the office manager. The conduct is exercised at the license compliance review (section 15.6), which is the gate through which the office's products are verified for license compliance.

A concrete example illustrates the disciplines' interaction. The resource economics specialist requests a high-resolution optical imagery acquisition for a commercial minerals commission. The operator receives the tasking request, verifies the request against the acquisition schedule, and confirms that the commercial charge code is applied. The operator tasks the imagery through the commercial provider, receives the imagery, and ingests it with the file integrity check, the format verification, and the metadata extraction. The operator records the license terms: the commercial license, the permitted use (commercial products only), the prohibited uses (no redistribution, no public-service products), and the attribution requirements. The operator configures the processing pipeline with the standard parameters for the data type and the intended domain, submits the processing jobs, and monitors the execution. The four disciplines are exercised in a single workflow:

the ingestion discipline (the acquisition was received and ingested correctly), the processing discipline (the standard parameters were applied), the base layer discipline (the base layers were verified for the project area), and the license discipline (the license terms were recorded and the commercial charge code was applied).

The four disciplines are exercised through the operator's daily rhythm and the office's procedures. The specialists receive the analysis-ready data and the current base layers they need for their model runs. The office manager authorizes the production on the basis of the technical lead's validation, which depends on the operator's processing. The national audit authority examines the license records and the base layer integrity records during the annual audit. The disciplines are established by the role family's training pipeline and qualifications framework, treated in Chapter 4, and maintained by the daily practice of the operator's work.

15.8 Worked Example: A Morning Ingestion Check with Multiple Issues

The data pipeline operator arrives at the office at 07:00 on a Tuesday morning. The operator opens the acquisition desk dashboard and begins the morning ingestion check (section 15.3). The overnight acquisition log shows three scheduled satellite acquisitions and two sensor feed subscriptions that should have delivered data overnight.

The first scheduled acquisition is a Sentinel-2 image over the agricultural region, expected at 06:00 UTC. The ingestion log shows that the image was received at 06:15 UTC (15 minutes after the expected time, within the normal delivery window) and ingested successfully. The file integrity check passed (the checksum matches the Copernicus Open Access Hub's checksum). The format verification passed (the Sentinel-SAFE format was correctly unpacked). The metadata extraction passed (the acquisition ID, the sensing time, the cloud cover percentage (15 per cent), and the tile identifier were extracted). The operator records the license terms (Copernicus free and open data, commercial use permitted with attribution) and releases the data to the processing queue with the public-service charge code PUB-2024-120. The provenance record is complete.

The second scheduled acquisition is a Landsat 8 image over the flood basin, expected at 04:30 UTC. The ingestion log shows that the image was received at 04:45 UTC but the file integrity check failed: the checksum does not match the EarthExplorer checksum. The operator investigates and discovers that the download was interrupted (the network connection was lost during the download, and the file was partially downloaded). The operator re-downloads the image from EarthExplorer, the second download's file integrity check passes, and the image is ingested successfully. The operator records the license terms (Landsat free and open data, no restrictions on use) and releases the data to the processing queue with the public-service charge code PUB-2024-121. The operator notes the initial download failure in the ingestion log, with the cause (network interruption) and the corrective action (re-download).

The third scheduled acquisition is a Maxar WorldView-3 image tasked for the energy infrastructure monitoring cycle, expected at 03:15 UTC. The ingestion log shows that the image was not received. The operator checks the Maxar tasking portal and discovers that the acquisition was declined by the provider: the cloud cover at the time of acquisition was 85 per cent, exceeding the 10 per cent maximum acceptable cloud cover specified in the tasking request. The operator notifies the technical lead: the WorldView-3 acquisition was declined due to cloud cover, and the energy

infrastructure monitoring cycle will be delayed until the next available acquisition window (48-72 hours, subject to cloud cover). The technical lead notifies the hazard specialist, who adjusts the monitoring cycle's schedule. The operator records the declined acquisition in the ingestion log, with the cause (cloud cover exceeded the threshold) and the next available window.

The first sensor feed subscription is the rain gauge network (15 gauges, 15-minute interval). The operator checks the latest data timestamp for each gauge: 14 gauges show data within the last 15 minutes (current), but Gauge RG-11 shows a last update timestamp of 03:00 UTC (over 4 hours ago). The operator checks the national hydrological service's status page and discovers that Gauge RG-11 is offline due to a communication link failure. The operator notifies the technical lead: Gauge RG-11 is offline, and the flood risk domain's rainfall data for the RG-11 area is stale. The technical lead notes the outage and determines that the flood model can use the remaining 14 gauges with interpolated rainfall for the RG-11 area. The operator records the sensor feed outage in the ingestion log, with the cause (communication link failure) and the impact (stale rainfall data for the RG-11 area).

The second sensor feed subscription is the river stage network (8 gauges, 15-minute interval). The operator checks the latest data timestamp for each gauge: all 8 gauges show data within the last 15 minutes. The operator checks the river stage values: all gauges are within normal range (below the 2-year flood threshold). No flood event is indicated.

The operator compiles the morning ingestion check report. The report covers: (1) Sentinel-2 image received and ingested successfully, released to the processing queue (PUB-2024-120); (2) Landsat 8 image initially failed (network interruption), re-downloaded and ingested successfully, released to the processing queue (PUB-2024-121); (3) Maxar WorldView-3 acquisition declined (cloud cover), energy infrastructure monitoring cycle delayed, technical lead and hazard specialist notified; (4) Rain gauge RG-11 offline (communication link failure), flood model to use interpolated data, technical lead notified; (5) River stage network fully operational, all gauges within normal range. The operator submits the report to the technical lead and proceeds to the processing pipeline configuration (section 15.4) with the two successfully ingested datasets.

This worked example illustrates the operator's four disciplines and the operational reality of the morning ingestion check. The ingestion discipline was exercised: two of three acquisitions were ingested, one was declined, and the cause and impact were documented. The processing discipline was exercised: the two ingested datasets were released to the processing queue with the correct charge codes and provenance records. The base layer discipline was exercised: the sensor feed outage (RG-11) was identified and reported, and the technical lead determined the mitigation (interpolated data). The license discipline was exercised: the license terms were recorded for both ingested datasets (Copernicus free and open, Landsat free and open). The operator's morning, which began with a routine check, encountered three issues (a network interruption, a cloud cover decline, and a sensor feed outage) and resolved or reported each one within the 30-minute check window, and the pipeline's operational reliability was maintained.

Chapter 16 — The Production Operator

16.1 The Role of the Production Operator

The production operator is the geospatial office's cartographic production and archival operator. The operator is accountable for the product production, the classification marking, the archival, and the access logging. The operator is the primary operator of the production and archive submodules of the geospatial information system, treated in Chapters 11 and 12 and in Chapter 29. The operator reports to the technical lead, who configures the production templates and the archive parameters, and performance is assessed by the office manager on the production function's quality and timeliness.

The operator's work is the back end of the pipeline. Every product the office delivers passes through the operator's hands at the production and archive stages. The quality of the operator's work determines whether the product is clear, correct, and usable — clear in its cartographic design, correct in its classification marking, and usable in its format and delivery. The operator's work is the last operational step before the product reaches the consumer, and any defect that has propagated through the pipeline emerges in the operator's hands.

The operator's role is shaped by the production and archive capabilities of the pipeline. The production capability transforms the specialists' model outputs into final deliverables through the cartographic design, the symbology application, the classification marking, and the format conversion. The archive capability preserves the office's outputs and maintains the catalog that makes them retrievable. The operator operates both capabilities, and the operator's work is the last operational step before the product reaches the consumer.

The operator's position is established at the operational tier of the role family, treated in Chapter 4. The position is held by an officer who has completed the training pipeline, holds the operator certification, and has served a probationary period under the supervision of a senior operator. The position is a career position, and the operator may progress to the senior operator rank and then to the technical lead position with the additional certifications and experience specified by the qualifications framework. The operator's training pipeline includes the foundational course (covering the pipeline architecture, the information system, the classification framework, and the office's procedures), the product studio course (covering the cartographic templates, the symbology library, the classification marking tools, and the format conversion tools), and the catalog and archive course (covering the product catalog, the version history, the retention and disposal, the access control, and the audit jurisdiction). The operator must pass the certification assessment for each course before operating the corresponding submodule independently. The certification is maintained through continuing education, with refresher courses and update courses required on a periodic basis to keep the operator's knowledge current with the evolving tools and methods. The operator also receives domain-specific familiarization training, which provides a basic understanding of each domain's products, classification regimes, and consumer requirements, so that the operator can apply the correct template and classification marking for each domain's products.

National mapping agencies organize the production and archive function in comparable ways. The United States Geological Survey's publication production specialists operate the cartographic production and archive pipeline for the Survey's map and report series, under the supervision of the center's production management team. The Ordnance Survey's data production specialists operate the

production and archive pipeline for the Survey's geospatial products, under the supervision of the data management team. IGN France's cartographic production specialists operate the production and archive pipeline for the Institute's map and data products, under the supervision of the production directorate. The geospatial office's production operator operates the production and archive pipelines under the supervision of the technical lead, with the production, classification, and archival conducted to the technical lead's configuration and the products reviewed by the office manager before release.

The operator's operational responsibilities are specified in the job breakdown sheets in sections 16.3 through 16.6. The product production run (section 16.3) transforms a specialist's model output into a final deliverable. The classification marking verification (section 16.4) verifies the correct classification marking before release. The archival and cataloguing (section 16.5) archives every product and maintains the catalog. The classified access log review (section 16.6) reviews the access logs for classified products. The operator conducts the operational steps of each procedure, and the technical lead configures the templates and the parameters while the office manager reviews the products before release.

16.2 The Operator's Daily Rhythm

The production operator's day is organized around the office's production on demand and the daily archival cycle. The day begins with the morning archival check, which verifies that the overnight archival was conducted correctly. The check is conducted at the catalog and archive submodule and takes approximately fifteen minutes.

The product production runs (section 16.3) are conducted as the specialists' model outputs become available. A production run is conducted for every product the office produces, and it comprises the application of the cartographic template, the symbology, the classification marking, and the format conversion. The run is conducted under the technical lead's configuration and the office manager's review. The operator who conducts the run well produces products that are clear, correct, and usable.

The classification marking verification (section 16.4) is conducted for every product before release. The verification verifies that every product carries the correct classification marking in the standard format, in the standard location, and consistent across all components. The verification is the last check before the product reaches the consumer, and it is conducted for every product without exception.

The archival and cataloguing (section 16.5) is conducted for every product after release. The archival preserves the released product in the catalog and archive submodule, with the standard metadata, the provenance record, the retention designation, and the access permissions. The cataloguing makes the product retrievable for re-use, for audit, and for training.

The classified access log review (section 16.6) is conducted at the end of each day. The review verifies that every access to a classified product was authorized, logged, and consistent with the access permissions. The review is the mechanism that detects unauthorized access promptly, and it is conducted at the catalog and archive submodule.

The operator's rhythm is distinct from the data pipeline operator's rhythm. The data pipeline operator operates the front of the pipeline — the acquisition and processing — and the production

operator operates the back of the pipeline — the production and archive. The two operators are the bookends of the pipeline, and the specialists' modeling work sits between them. The two operators' rhythms are interdependent: the production operator cannot produce a product until the data pipeline operator has provided the processed data, and the data pipeline operator's work is incomplete until the production operator has archived the product.

The daily rhythm is the operator's mechanism for maintaining the production and archive functions' reliability. An operator who skips the morning archival check does not know whether the overnight archival was conducted correctly. An operator who skips the classification marking verification does not know whether the product is correctly marked. An operator who skips the classified access log review does not know whether the day's access was authorized. The discipline is the daily practice of being present at the production and archive submodules and informed about the production function's state.

The operator's weekly rhythm includes the archive integrity audit, conducted on the first working day of each week, which verifies that the archive's checksums match the released products' checksums and that the catalog metadata is complete. The monthly rhythm includes the retention review, conducted on the first working day of each month, which identifies the products whose retention period has expired and prepares them for secure disposal in accordance with the records management standards. The operator's weekly and monthly rhythms complement the daily rhythm and ensure that the production and archive functions are maintained over time.

16.3 The Product Production Run

The product production run is the production operator's core task of transforming a specialist's model output into a final deliverable. The run is conducted for every product the office produces, and it comprises the application of the cartographic template, the symbology, the classification marking, and the format conversion. The run is conducted under the technical lead's configuration and the office manager's review.

Important Steps	Key Points	Reasons for Key Points
1. Receive the model output from the specialist and verify its completeness.	Confirm that the model output includes the analytical result, the validation record, and the specialist's authorization for production.	An incomplete model output must not enter the production run, and the specialist's authorization is the signal that the analytical work is finished.
2. Select the cartographic template appropriate for the product type.	Choose the template that matches the product type (flood inundation map, crop condition bulletin, infrastructure risk map, prospectivity study) and the classification level.	The template determines the cartographic design, and the wrong template produces a product that does not meet the office's cartographic standards.

Important Steps	Key Points	Reasons for Key Points
3. Apply the symbology to the model output.	Apply the symbology from the cartographic standards, and verify that the symbology is consistent with the office's other products of the same type.	Consistent symbology allows the consumer to interpret the product without relearning the symbology, and inconsistency undermines the office's credibility.
4. Apply the classification marking to the product.	Determine the classification level from the product's content and the domain's classification regime, and apply the marking in the standard format and location.	The classification marking is the security integrity guarantee, and an incorrect marking is a security breach that could result in the disclosure of classified information.
5. Convert the product to the formats required by the consumer.	Produce the geospatial format for consumers who ingest the product into their systems, and the report format for consumers who read the product.	The format conversion is the usability guarantee, and a product in the wrong format is a product that the consumer cannot use.
6. Submit the draft product to the office manager for review.	Submit the draft with the model output, the template, the symbology, the classification marking, and the format conversion, and note any issues that arose during the production run.	The office manager's review is the quality gate of the production function, and the manager cannot review what is not submitted.
7. Release the approved product to the consumer through the appropriate distribution channel.	For public-service products, release through the public-service distribution channel; for commercial products, release through the commercial channel only after the office manager's commercial release authorization.	The distribution channel separation is the distribution mechanism of the firewall, and the wrong channel is a firewall breach.

16.4 The Classification Marking Verification

The classification marking verification is the production operator's task of verifying that every product carries the correct classification marking before the product is released. The verification is conducted for every product, without exception, and it is the last check before the product reaches the consumer. The verification prevents the release of misclassified products and the security breaches that follow.

Important Steps	Key Points	Reasons for Key Points
1. Determine the correct classification level for the product.	Determine the level from the product's content and the domain's classification regime: Public, Internal, Confidential, Medical Confidential, or Security Confidential.	The classification level is determined by the content's sensitivity, and the wrong level is a security breach that could result in disclosure or in improper restriction.
2. Verify that the classification marking is applied in the standard format.	Check that the marking is in the standard format, with the classification level, the product identifier, and the handling caveats (where applicable).	The standard format allows the consumer to recognize the classification immediately, and non-standard formats may be overlooked or misinterpreted.
3. Verify that the classification marking is in the standard location.	Check that the marking is in the standard location on the product: the header and footer for documents, the title block for maps, and the metadata for geospatial products.	The standard location allows the consumer to find the marking immediately, and markings in non-standard locations may be overlooked.
4. Verify that the classification marking is consistent across all components of the product.	Check that the marking is the same on the cover, the body, the appendices, and the metadata, and that there are no components with missing or inconsistent markings.	Inconsistent markings create confusion about the product's classification, and the confusion is a security risk.
5. Verify that the product's content matches its classification marking.	Review the product's content for any information that is more sensitive than the marking indicates, and escalate any discrepancy to the office manager.	A product that contains information more sensitive than its marking is a security breach, and the breach must be corrected before the product is released.

Important Steps	Key Points	Reasons for Key Points
6. For commercial products, verify the Confidential marking includes the customer identity.	Confirm that the marking includes the customer identity and the product identifier, and that the marking restricts the product to the customer and the office.	The commercial marking is the mechanism that enforces the firewall's classification mechanism, and the customer identity is the restriction that prevents disclosure to non-minerals staff.
7. Record the verification in the production log.	Record the product identifier, the classification level, the verification result, and the operator's name in the production log.	The production log is the compliance record for the classification marking verification, and it is the record that the national audit authority examines during the annual audit.

16.5 The Archival and Cataloguing

The archival and cataloguing is the production operator's task of archiving every product the office produces and maintaining the catalog that makes the products retrievable. The archival is conducted for every product, without exception, and it is the mechanism that preserves the office's institutional memory. The archival is conducted at the catalog and archive submodule.

Important Steps	Key Points	Reasons for Key Points
1. Archive a copy of every product in the catalog and archive submodule.	Archive the final product in all produced formats, and verify that the archive copy is identical to the released copy.	The archive is the permanent record of the office's work, and the archive copy must be the product that was released, not a draft or a variant.
2. Catalogue the product with the standard metadata.	Record the product identifier, the product type, the production date, the classification marking, the retention designation, the provenance record, and the access permissions.	The catalog is the index that makes the product retrievable, and incomplete metadata makes the product unfindable.

Important Steps	Key Points	Reasons for Key Points
3. Link the product's provenance record to the acquisition, processing, and modeling provenance.	Verify that the provenance chain is complete from the source data through the processing, the modeling, and the production to the archive.	The provenance chain is the integrity guarantee of the pipeline, and a break in the chain at the archive stage is a defect that must be corrected.
4. Assign the retention designation to the product.	Assign the retention period specified by the records management standards for the product type, and record the retention designation in the catalog.	The retention designation determines how long the product is retained, and the designation must be assigned at the moment of archival to prevent premature disposal.
5. Set the access permissions for the product.	Set the access permissions according to the product's classification marking: Public for public products, Internal for internal products, and restricted for Confidential and Security Confidential products.	The access permissions are the mechanism that enforces the classification marking, and the wrong permissions could result in unauthorized access.
6. Verify the archival integrity.	Run the archival integrity check, which verifies that the archived product's checksum matches the released product's checksum, and that the metadata is complete.	The integrity check ensures that the archived product is identical to the released product, and that the metadata is sufficient for retrieval and audit.
7. Record the archival in the catalog and archive submodule.	Record the product identifier, the archive date, the retention designation, the access permissions, and the archival integrity check result.	The archival record is the compliance record for the archive function, and it is the record that the national audit authority examines during the annual audit.

16.6 The Classified Access Log Review

The classified access log review is the production operator's daily task of reviewing the access logs for the office's classified products. The review is conducted at the end of each day, and it verifies that every access to a classified product was authorized, logged, and consistent with the access permissions. The review detects unauthorized access to classified products and supports the office manager's investigation of any breach.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the day's access logs from the catalog and archive submodule.	Retrieve the logs for all classified products: Confidential, Medical Confidential, and Security Confidential.	The access logs are the record of who accessed what, and the logs must be reviewed daily to detect unauthorized access promptly.
2. Verify that each access was by an authorized person.	Cross-check each access against the access permissions for the product, and flag any access by an unauthorized person.	Unauthorized access is a security breach, and the breach must be detected and reported promptly to prevent further access.
3. Verify that each access was for an authorized purpose.	Review the access purpose recorded by the accessor, and flag any access that appears to be for a purpose inconsistent with the office's work.	Authorized persons may access products for unauthorized purposes, and the purpose review is the check that detects such access.
4. Flag any unusual access patterns.	Flag any access that occurs outside normal working hours, any access by a person who does not normally access the product type, and any bulk access that appears to be data exfiltration.	Unusual access patterns may indicate a security incident, and the patterns must be flagged for investigation even if each individual access appears authorized.
5. Report any unauthorized access or unusual patterns to the office manager.	Provide a clear summary of the issue, the affected products, the persons involved, and the recommended corrective action.	The office manager is the officer who can authorize the corrective action, and the operator's summary is the information the manager needs to act.

Important Steps	Key Points	Reasons for Key Points
6. Record the review results in the catalog and archive submodule.	Record the logs reviewed, the access verified, the issues flagged, and the corrective actions taken or recommended.	The review record is the compliance record for the access control function, and it is the record that the national audit authority examines during the annual audit.

16.7 The Operator's Discipline

The production operator's discipline is the daily practice of maintaining the production and archive functions' integrity under operational pressure. The discipline has four components: production, classification, archival, and access. Each component is maintained through a specific procedure that surfaces the pressure and verifies the standard.

The production discipline covers the cartographic standards: the cartographic templates, the symbology, the generalization rules, and the format specifications, which are configured by the technical lead and applied by the operator. Three pressures bear on it. Volume pressure arises when the production queue is full and the operator is tempted to skip the cartographic review. Consumer urgency pressure arises when the consumer needs the product urgently and the operator is tempted to release it without the full production cycle. Consistency pressure arises when the operator is tempted to deviate from the standard symbology to satisfy a consumer preference. The operator's production discipline is the application of the cartographic standards to every product. The application is exercised at the product production run (section 16.3), which is the gate through which the specialist's model output becomes a delivered product.

The classification discipline covers the application of the classification level that the content's sensitivity requires, in the standard format, in the standard location, and consistent across all components. Three pressures bear on it. Handling burden pressure arises when a higher classification marking requires more handling, which slows the release. Consumer access pressure arises when the consumer needs the product but does not have the clearance for the higher marking. Default pressure arises when the lowest classification level is the easiest to apply. The operator's classification discipline is the application of the correct classification marking to every product. The application is exercised at the classification marking verification (section 16.4), which is the gate through which the product is verified for release.

The archival discipline covers the archiving of every product the office produces, with the standard metadata, the provenance record, the retention designation, and the access permissions, and with the integrity verification that the archived product's checksum matches the released product's checksum. Three pressures bear on it. Time pressure arises when the archival takes time and the operator is tempted to defer it. Metadata pressure arises when the metadata is detailed and the operator is tempted to omit fields. Volume pressure arises when the production volume is high and the operator is tempted to batch the archival and risk omission. The operator's archival discipline is the archiving of every product at the moment of release, with the complete metadata and the integrity

verification. The conduct is exercised at the archival and cataloguing (section 16.5), which is the gate through which the released product enters the archive.

The access discipline covers the assignment of the access permissions according to the product's classification marking, the logging of every access to a classified product, and the daily review of the access logs. Three pressures bear on it. Volume pressure arises when the access logs are many and the review takes time. Familiarity pressure arises when the accessors are familiar colleagues and the operator is tempted to assume the access was authorized. Urgency pressure arises when the consumer needs the product urgently and the operator is tempted to grant access without verifying the authorization. The operator's access discipline is the review of every access. The conduct is exercised at the classified access log review (section 16.6), which is the gate through which the day's access is verified.

A concrete example illustrates the disciplines' interaction. The hazard specialist submits a flood inundation map for a developing event, with the Internal classification marking and a delivery deadline of two hours. The operator receives the model output, verifies the specialist's authorization, and selects the flood inundation map template. The operator applies the symbology (the blue depth gradient), the classification marking (Internal, in the header and footer and the title block), and the format conversion (GeoTIFF for the emergency management authority's operational system and PDF for the briefing pack). The operator submits the draft to the office manager, who reviews the product and authorizes the release through the public-service distribution channel. The operator archives the product with the standard metadata, the provenance record, the retention designation (permanent for flood inundation maps), and the access permissions (Internal). The four disciplines are exercised in a single production cycle: the production discipline (the cartographic template and symbology were applied), the classification discipline (the Internal marking was verified), the archival discipline (the product was archived with complete metadata), and the access discipline (the access permissions were set for the Internal marking).

The four disciplines are exercised through the operator's daily rhythm and the office's procedures. The consumers receive the products. The emergency management authority receives flood intelligence that is clearly presented and correctly classified. The strategic reserve authority receives the weekly crop condition bulletin in the format the authority's systems can ingest. The commercial customers receive prospectivity products that are correctly marked and released through the commercial channel. The national audit authority examines the archive's completeness, the classification compliance, and the access logs during the annual audit. The disciplines are established by the role family's training pipeline and qualifications framework, treated in Chapter 4, and maintained by the daily practice of the operator's work.

16.8 Worked Example: Producing and Archiving a Commercial Prospectivity Report

The resource economics specialist has completed a mineral prospectivity assessment for Meridian Resources, a mining company that has commissioned a custom prospectivity map for a 500-square-kilometre area in the northern highlands (commission COMM-2024-017, agreed price $45,000, delivery timeline 8 weeks). The specialist has produced the prospectivity map using the weights-of-evidence method (validated against known deposits, 78 per cent prospectivity accuracy,

prediction ratio 4.2), has documented the model run in the provenance record (prospectivity plug-in version 3.1, weights-of-evidence method, training data from the geological survey's mineral occurrence database), and has authorized the model output for production.

The production operator receives the model output from the modeling workbench. The operator verifies the completeness: the analytical result (the prospectivity map, a raster at 30-metre resolution with the prospectivity scores from 0 to 1), the validation record (78 per cent accuracy, 4.2 prediction ratio, both within tolerance), and the specialist's authorization for production are all present. The operator proceeds to the production run (section 16.3).

The operator selects the cartographic template for the mineral prospectivity map with Confidential classification (because the product is a commercial commission for Meridian Resources, and the customer identity and the specific findings are commercially sensitive). The template specifies the UTM Zone 42N projection, the A3 landscape page layout, the legend in the lower right corner, and the classification marking at the top and bottom of every page.

The operator applies the symbology. The prospectivity raster is symbolized with the warm-cool gradient: the highest prospectivity (scores 0.8 to 1.0) in dark red (CMYK 0/100/100/0), high prospectivity (0.6 to 0.8) in orange (CMYK 0/50/100/0), moderate prospectivity (0.4 to 0.6) in yellow (CMYK 0/15/100/0), low prospectivity (0.2 to 0.4) in light blue (CMYK 40/0/0/0), and the lowest prospectivity (0.0 to 0.2) in dark blue (CMYK 80/0/0/20). The geological formations are overlaid in muted browns and greens at 30 per cent transparency, so that the geology is visible but does not dominate the prospectivity. The known deposits are marked as black stars (10-point), and the exploration licenses are outlined in black dashed lines (1-point). The operator verifies that the symbology is consistent with the office's previous prospectivity maps.

The operator applies the classification marking. The marking is "CONFIDENTIAL" in capital letters, centred at the top and bottom of every page, with the customer identity ("Meridian Resources") and the product identifier ("COMM-2024-017") included in the marking. The marking is also embedded in the GeoTIFF's metadata tags and the PDF's document properties. The operator verifies the marking using the classification marking verification procedure (section 16.4): the marking is in the standard format, in the standard location (header, footer, title block, metadata), consistent across all components, and the content matches the marking (no information more sensitive than Confidential is present). The operator records the verification in the production log.

The operator conducts the format conversion. Meridian Resources has requested the product in two formats: GeoTIFF (for ingestion into the company's GIS system, which uses MapInfo Pro) and PDF (for the commission report, which the company's geologists will review). The operator converts the prospectivity raster to GeoTIFF (with the georeferencing in UTM Zone 42N, EPSG:32642), the geological formations and the known deposits to Shapefile (with the .prj, .shp, .shx, and .dbf files), and the full product (the map layout with the symbology, the legend, and the classification marking) to PDF. The operator verifies the format conversion: the GeoTIFF's georeferencing is correct, the Shapefile's projection file is present, and the PDF's classification marking is applied.

The operator submits the draft product to the office manager for review. The office manager reviews the product under the commercial release authorization procedure (section 13.4): the firewall compliance (the product was produced from commercial data sources and commercial model runs, verified through the charge code COMM-2024-017), the classification marking (CONFIDENTIAL

with the customer identity and product identifier), the charge code (COMM-2024-017, the commercial charge code), the customer identity (Meridian Resources, verified against the commercial customer register), and the agreed price ($45,000, verified against the pricing schedule). The manager signs the commercial release log and directs the operator to release the product through the commercial distribution channel.

The operator releases the product to Meridian Resources through the commercial distribution channel (the secure file transfer protocol, with the download link sent to the Meridian Resources contact person, and the link expiring after 30 days). The release is recorded in the catalog and archive submodule: the product identifier (MINE-2024-017-v01), the release date, the customer, the distribution channel, the formats (GeoTIFF, Shapefile, PDF), the classification marking (CONFIDENTIAL), and the office manager's commercial release authorization.

The operator archives the product (section 16.5). The operator archives a copy of the product in all three formats in the catalog and archive submodule. The operator records the catalog entry: product identifier (MINE-2024-017-v01), product type (mineral prospectivity map), production date, classification marking (CONFIDENTIAL), domain (minerals), retention designation (10 years, per the commission contract), access permissions (restricted to the office manager and the resource economics specialist). The operator links the provenance record to the model run (prospectivity plug-in version 3.1, weights-of-evidence, the training data from the mineral occurrence database) and to the acquisition (the commercial geological data, the commercial aeromagnetic data, the commercial geochemical data, all acquired under the commercial charge code COMM-2024-017). The operator runs the archival integrity check (the archived product's checksum matches the released product's checksum) and records the archival in the catalog.

At the end of the day, the operator conducts the classified access log review (section 16.6). The review covers the day's access to the CONFIDENTIAL products in the archive. The access log shows three accesses to MINE-2024-017-v01: one by the production operator (for the archival), one by the resource economics specialist (for a final review of the archived product), and one by the office manager (for the commercial release log verification). All three accesses are by authorized personnel, all three are for authorized purposes, and no unusual access patterns are identified. The operator records the review results in the catalog and archive submodule.

This worked example illustrates the production operator's four disciplines operating together in a commercial product context. The production discipline was exercised (the cartographic template, the symbology, and the format conversion were applied). The classification discipline was exercised (the CONFIDENTIAL marking with the customer identity was applied and verified). The archival discipline was exercised (the product was archived with complete metadata, provenance, and retention designation). The access discipline was exercised (the access permissions were set to restricted, and the day's access was reviewed and verified). The four disciplines, exercised through the operator's daily rhythm and the office's procedures, produced a commercial product that was professionally produced, correctly classified, properly archived, and properly access-controlled, and the product's release through the commercial channel maintained the firewall between the commercial and public-service arms.

Chapter 17 — The Hazard Specialist

17.1 The Role of the Hazard Specialist

The hazard specialist is the geospatial office's specialist for the flood risk and energy infrastructure domains. The specialist is accountable for the hydrological modeling, the change detection, the event response, and the validation. The specialist's domain track is the public-service, non-commercial side of the office. The specialist operates the hydrology and change detection plug-ins of the modeling workbench, treated in Chapter 10. The specialist reports to the technical lead, who configures the modeling workbench and validates the specialist's model runs, and performance is assessed by the office manager on the quality and timeliness of the specialist's products.

The specialist's work produces the intelligence that the emergency management authority and the energy sector entity consume. The flood risk domain produces the flood inundation maps, the exposure reports, and the evacuation guidance that the emergency management authority uses for the operational response. The energy infrastructure domain produces the infrastructure risk maps, the encroachment alerts, and the maintenance advisories that the energy sector entity and the sectoral regulator use for the asset management and the regulatory compliance. The two domains are distinct in their cadences, their consumers, and their classification regimes.

The flood risk domain operates on a continuous cadence with event-driven production. The specialist monitors the sensor feeds continuously, conducts the flood model run on the standard schedule (daily during the flood season, weekly otherwise), and responds to flood events with event-driven production that overrides the scheduled rhythm. The consumer is the emergency management authority, and the classification regime is tiered: Public for the situational awareness products, Internal for the operational planning products, and Security Confidential for the products that identify critical infrastructure at risk.

The energy infrastructure domain operates on a monthly cadence with event-driven response. The specialist conducts the change detection run on a monthly cycle, monitoring the high-resolution imagery for changes that affect the energy infrastructure. The consumers are the energy sector entity and the sectoral regulator, and the classification regime is Security Confidential by default: the geolocation of critical energy infrastructure is security-sensitive, and the products that identify the infrastructure and its vulnerabilities are restricted to the security-cleared personnel.

The specialist who serves both domains must switch between the cadences, the consumers, and the classification regimes without confusion. The switching is the discipline that the specialist's training and the office's procedures are designed to support. The flood model run (section 17.3) is conducted on the continuous cadence with the tiered classification regime. The energy infrastructure monthly monitoring cycle (section 17.4) is conducted on the monthly cadence with the Security Confidential classification regime. The flood event response (section 17.5) overrides both cadences for the duration of the event. The change detection validation (section 17.6) is conducted on every change detection run.

The specialist's position is established at the operational tier of the role family, treated in Chapter 4. The position is held by an officer who has completed the training pipeline, holds the specialist certification in the hazard domain, and has served a probationary period under the supervision of a senior specialist. The position is a career position, and the specialist may progress to the senior

specialist rank and then to the technical lead position with the additional certifications and experience specified by the qualifications framework. The specialist's training pipeline includes the foundational course (covering the pipeline architecture, the information system, the classification framework, and the office's procedures), the modeling workbench course (covering the workbench's data interface, parameter management, validation framework, and provenance recording), the hydrology plug-in course (covering the HEC-HMS rainfall-runoff model, the HEC-RAS 2D hydraulic model, the inundation mapping, and the validation against observed events), and the change detection plug-in course (covering the image-to-image and object-to-object change detection algorithms, the encroachment detection, the vulnerability scoring, and the validation against ground truth inspections). The specialist must pass the certification assessment for each course before operating the corresponding plug-in independently. The certification is maintained through continuing education, with refresher courses and update courses required on a periodic basis. The specialist also receives domain-specific continuing education through attendance at professional conferences (such as the American Geophysical Union Fall Meeting for hydrology, or the IEEE International Geoscience and Remote Sensing Symposium for remote sensing), which the office supports through conference fee reimbursement and travel funding.

National mapping agencies organize the hazard and infrastructure specialist function in comparable ways. The United States Geological Survey's hydrologists and remote sensing specialists conduct the flood modeling and the infrastructure monitoring for the Survey's hazard programs, in coordination with the Federal Emergency Management Agency and the Department of Energy. The European Commission's Joint Research Centre conducts the European Flood Awareness System and the Copernicus Emergency Management Service, with hydrologists and remote sensing specialists conducting the modeling and the monitoring. Geoscience Australia's hazard specialists conduct the flood modeling and the infrastructure monitoring for the agency's hazard programs, in coordination with the state emergency services and the energy regulators. The geospatial office's hazard specialist conducts the flood modeling and the infrastructure monitoring under the supervision of the technical lead, with the modeling and the monitoring conducted to the technical lead's configuration and the model runs validated against independent ground truth before release.

17.2 The Specialist's Daily Rhythm

The hazard specialist's day is organized around the variable modeling cadence of the two domains. The day begins with the flood risk monitoring, which is conducted continuously through the day as the sensor feeds update. The monitoring comprises the review of the rain gauge and river stage alerts, the review of the weather forecast, and the assessment of the flood risk for the monitored basins. The monitoring is brief on a quiet day — fifteen to thirty minutes — and extended on a day with active alerts or developing events.

The energy infrastructure monitoring is conducted on the monthly cycle, with the specialist conducting the change detection run on the first working day of each month. The monitoring comprises the retrieval of the latest high-resolution imagery, the configuration of the change detection plug-in, the execution of the change detection run, the review of the detected changes, the cross-referencing with the infrastructure asset registry, and the generation of the infrastructure risk

map, the encroachment alerts, and the maintenance advisories. The monitoring takes approximately two days per month.

The validation of the specialist's model runs is conducted on every run. The flood model run is validated against the observed flood events, with the inundation extent and depth compared with the ground truth observations. The change detection run is validated against the ground truth inspections provided by the energy sector entity, with the detected changes compared with the inspection reports. The validation is the analytical integrity guarantee of the specialist's work, and a run that fails validation is not released.

The event response, when activated by the office manager under the incident response activation procedure (section 13.6), overrides the scheduled rhythm. The specialist's role in the response is to produce the flood intelligence that the emergency management authority needs for the operational response. The response is conducted under time pressure, and the specialist's discipline is the practice of maintaining the quality standards despite the pressure.

The specialist's rhythm is distinct from the operators' rhythms. The operators operate the pipeline on a fixed cadence; the specialist operates the models on a variable cadence. The specialist pulls from the pipeline on the specialist's own schedule: the flood risk monitoring is continuous, the energy infrastructure monitoring is monthly, and the event response is event-driven. The specialist's rhythm produces the office's analytical intelligence, and the timeliness and the quality of the specialist's products depend on the rhythm.

The specialist's weekly rhythm includes the flood model calibration check, conducted on the first working day of each week, which verifies that the hydrological model's parameters (the Manning's n roughness coefficients, the SCS curve numbers, the baseflow recession constants) are still appropriate for the current season and the current catchment conditions. The calibration check is brief on a quiet week, but it is the mechanism by which the specialist detects the gradual changes in the catchment (the vegetation growth that changes the roughness, the urban expansion that changes the curve numbers, the groundwater depletion that changes the baseflow) that would, if undetected, degrade the model's accuracy over time. The specialist's monthly rhythm includes the energy infrastructure monitoring cycle, conducted on the first working day of each month, and the model version review, conducted in coordination with the technical lead, which verifies that the authorized model versions are still the best available and that no new versions have been released by the developers that warrant evaluation. The daily rhythm is the specialist's mechanism for maintaining the analytical integrity of the two domains. A specialist who does not monitor the flood risk continuously misses the developing events. A specialist who does not conduct the change detection run on the monthly cycle misses the infrastructure changes. A specialist who does not validate every run produces intelligence that may be wrong in ways that are not visible to the consumer. The discipline is the daily practice of being present at the modeling workbench and informed about the two domains' states.

17.3 The Flood Model Run

The flood model run is the hazard specialist's core task for the flood risk domain, conducted on a continuous cadence with event-driven production. The run produces the flood inundation maps, the exposure reports, and the evacuation guidance that the emergency management authority consumes. The run is conducted at the modeling workbench using the hydrology plug-in, and it consumes the digital elevation model, the land cover, the hydrography, the rainfall data, and the river stage data. The run is conducted under the technical lead's configuration and validated against observed events.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the latest processed data from the processing engine.	Verify that the processed data includes the latest rainfall data, river stage data, and elevation data, and that the data's provenance record is complete.	The flood model's accuracy depends on the currency and the completeness of the input data, and the provenance record is the integrity guarantee that the data is what it claims to be.
2. Configure the hydrology plug-in for the model run.	Set the rainfall-runoff model parameters, the hydraulic model parameters, and the inundation model parameters per the standard parameter set, unless the technical lead has authorized non-standard parameters.	The model parameters determine the model's behavior, and non-standard parameters must be authorized to prevent the specialist from inadvertently producing non-standard products.
3. Execute the model run and monitor its execution.	Monitor the model's convergence, the runtime, and the memory usage, and note any anomalies that may indicate a numerical instability or a data quality issue.	Numerical instabilities and data quality issues can produce plausible but wrong results, and the anomalies are the early warning signs that the specialist must not ignore.
4. Validate the model run against observed data.	Compare the model's inundation extent and depth with the latest ground truth observations, and calculate the accuracy metrics per the domain's validation protocol.	The validation is the analytical integrity guarantee of the flood risk domain, and the model run that fails validation must not be released.

Important Steps	Key Points	Reasons for Key Points
5. Generate the flood inundation map, the exposure report, and the evacuation guidance.	Generate the products in the standard formats, with the classification marking determined by the content's sensitivity.	The products are the office's intelligence delivered to the emergency management authority, and the standard formats and the correct classification markings are the conditions of the products' usability and security.
6. Submit the products to the production operator for production.	Submit the products with the model run's provenance record, the validation results, and the specialist's authorization for production.	The production operator cannot produce the products without the specialist's authorization, and the provenance record and the validation results are the integrity guarantees that the production operator relies on.
7. Record the model run in the catalog and archive submodule.	Record the model version, the parameters, the run date, the validation results, and the products generated, linking the record to the processing and acquisition provenance.	The model provenance is the continuation of the provenance chain, and the record is the audit trail that allows the office to reconstruct the model run in the future.

17.4 The Energy Infrastructure Monthly Monitoring Cycle

The energy infrastructure monthly monitoring cycle is the hazard specialist's recurring task for the energy infrastructure domain, conducted on a monthly cadence with event-driven response. The cycle produces the infrastructure risk map, the encroachment alerts, and the maintenance advisories that the energy sector entity and the sectoral regulator consume. The cycle is conducted at the modeling workbench using the change detection plug-in, and it consumes the high-resolution imagery, the infrastructure asset registry, the elevation data, and the vegetation cover data. The cycle is conducted under the technical lead's configuration and validated against ground truth inspections.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the latest high-resolution imagery for the monitored infrastructure corridors.	Verify that the imagery was acquired within the monitoring cycle's window, and that the imagery's provenance record and license terms are recorded.	The change detection depends on the temporal consistency of the imagery, and imagery from outside the monitoring window may miss changes that occurred between acquisitions.
2. Configure the change detection plug-in for the monitoring run.	Set the change detection algorithm, the change threshold, the encroachment detection parameters, and the vulnerability scoring parameters per the standard parameter set.	The change detection parameters determine the sensitivity of the detection, and the wrong parameters produce either too many false alarms or missed detections.
3. Execute the change detection run and review the detected changes.	Review each detected change for plausibility, and distinguish between true changes (construction, vegetation growth, ground subsidence) and false alarms (shadow changes, seasonal variations, image artifacts).	The change detection algorithm produces both true and false detections, and the specialist's judgment is the filter that separates them, preventing false alarms from overwhelming the consumer.
4. Cross-reference the detected changes with the infrastructure asset registry.	Identify the assets affected by each change, and assess the vulnerability of each affected asset using the vulnerability scoring algorithm.	The vulnerability scoring is the mechanism that prioritizes the changes for the consumer, and the cross-referencing is the mechanism that connects the detected changes to the specific assets the consumer manages.

Important Steps	Key Points	Reasons for Key Points
5. Generate the infrastructure risk map, the encroachment alerts, and the maintenance advisories.	Generate the products with the Security Confidential classification marking, as the geolocation of critical energy infrastructure is security-sensitive by default.	The Security Confidential marking is the security integrity guarantee for the energy infrastructure domain, and the marking must be applied to every product unless the office manager has explicitly declassified the product.
6. Validate the monitoring run against ground truth inspections.	Compare the detected changes with the ground truth inspection reports provided by the energy sector entity, and calculate the accuracy metrics per the domain's validation protocol.	The validation is the analytical integrity guarantee of the energy infrastructure domain, and the validation against ground truth is the check that confirms the change detection's accuracy.
7. Submit the products to the production operator and record the run in the catalog and archive submodule.	Submit the products with the provenance record and the validation results, and record the run with the model version, parameters, and validation results.	The submission and the recording complete the monitoring cycle, and the provenance and the validation are the integrity guarantees that the production operator and the archive rely on.

17.5 The Flood Event Response

The flood event response is the hazard specialist's event-driven task, conducted when a flood event is imminent or occurring. The response overrides the scheduled rhythm and takes priority over all other work. The response is activated by the office manager under the incident response activation procedure (section 13.6), and the specialist's role is to produce the flood intelligence that the emergency management authority needs for the operational response. The response is conducted under time pressure.

Important Steps	Key Points	Reasons for Key Points
1. Receive the activation briefing from the office manager.	Confirm understanding of the event, the required products, the delivery timeline, and the classification level; ask for clarification on any point that is unclear.	The briefing is the specialist's direction for the response, and unclear direction produces the wrong intelligence, which is worse than no intelligence in an emergency.
2. Retrieve the latest sensor feed data and the latest processed imagery.	Verify that the sensor feeds are active and delivering current data, and that the processed imagery is the most recent available; flag any data gaps to the data pipeline operator immediately.	The flood event response depends on current data, and data gaps must be identified and addressed immediately to prevent the response from working against stale or incomplete data.
3. Execute the flood model run with the event parameters.	Use the event-specific parameters (the observed rainfall, the observed river stages, the forecast rainfall) rather than the standard parameters, and document the parameter changes.	The event-specific parameters reflect the actual conditions, and the standard parameters may not be appropriate for the event's severity; the documentation is the audit trail for the non-standard parameters.
4. Validate the model run against the observed event data.	Conduct the validation even under time pressure; do not waive the validation for speed, as an unvalidated model run may produce plausible but wrong inundation maps.	The validation is the analytical integrity guarantee, and waiving it under time pressure produces intelligence that may be wrong in ways that are not visible to the consumer, which is particularly dangerous in an emergency.

Important Steps	Key Points	Reasons for Key Points
5. Generate the event-specific flood inundation map, exposure report, and evacuation guidance.	Generate the products with the classification level specified in the activation briefing, and include the model run's timestamp and the data's cutoff time.	The timestamp and the cutoff time are critical for the emergency management authority's operational planning, as the authority needs to know how current the intelligence is.
6. Deliver the products to the office manager for review and release.	Deliver the products with the provenance record and the validation results, and brief the manager on any issues that arose during the response.	The manager's review is the quality gate, and the briefing is the mechanism that ensures the manager is aware of any issues that may affect the intelligence's reliability.
7. Continue the response until the event is resolved and the office manager stands down the response.	Produce updated intelligence as the event evolves, and document the response timeline, the products produced, and the issues encountered for the post-incident review.	The flood event evolves over time, and the intelligence must evolve with it; the documentation is the basis for the post-incident review and the office's learning from the event.

17.6 The Change Detection Validation

The change detection validation is the hazard specialist's quality assurance task for the energy infrastructure domain, conducted on every change detection run. The validation verifies that the detected changes are accurate, by checking the detected changes against ground truth inspections and against the specialist's expert judgment. A run that fails validation is not released.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the ground truth inspection reports from the energy sector entity.	Request the reports for the monitored infrastructure corridors, covering the same period as the change detection run.	The ground truth inspections are the independent verification of the change detection, and the reports must cover the same period for the comparison to be meaningful.

Important Steps	Key Points	Reasons for Key Points
2. Compare the detected changes with the ground truth inspections.	For each detected change, determine whether the ground truth confirms the change (true positive), contradicts the change (false positive), or is silent; for each ground truth change, determine whether the detection captured it (true positive) or missed it (false negative).	The comparison is the basis for the accuracy metrics, and the true positive, false positive, and false negative counts are the inputs to the metrics.
3. Calculate the accuracy metrics per the domain's validation protocol.	Calculate the detection rate (true positives divided by all actual changes), the false alarm rate (false positives divided by all detected changes), and the overall accuracy.	The accuracy metrics are the quantitative measure of the change detection's performance, and the metrics must be calculated consistently across runs for the performance to be comparable over time.
4. Review the accuracy metrics against the domain's tolerance.	If the metrics are within the tolerance, the run is validated; if the metrics are outside the tolerance, investigate the cause and either adjust the parameters and re-run or document the discrepancy and escalate to the technical lead.	The tolerance is the threshold that separates acceptable from unacceptable performance, and runs outside the tolerance must not be released without investigation and correction.
5. Record the validation results in the catalog and archive submodule.	Record the ground truth source, the comparison results, the accuracy metrics, the tolerance, and the validation outcome.	The validation record is the analytical integrity guarantee of the energy infrastructure domain, and it is the record that the national audit authority examines during the annual audit.

17.7 The Specialist's Discipline

The hazard specialist's discipline is the daily practice of maintaining the analytical integrity of the flood risk and energy infrastructure domains under operational pressure. The discipline has four

components: modeling, validation, event, and classification. Each component is maintained through a specific procedure that surfaces the pressure and verifies the standard.

The modeling discipline covers the model standards: the model parameter sets, the model version control, the model run documentation, and the model provenance record, which are configured by the technical lead and applied by the specialist. Three pressures bear on it. Parameter pressure arises when the standard parameters are wrong for the event and the specialist is tempted to deviate without authorization. Version pressure arises when the latest model version is untested and the specialist is tempted to use it without the validation. Documentation pressure arises when the documentation takes time and the specialist is tempted to omit it. The specialist's modeling discipline is the application of the standard parameters, the use of the authorized model versions, and the documentation of every run. The conduct is exercised at the flood model run (section 17.3) and the energy infrastructure monthly monitoring cycle (section 17.4).

The validation discipline covers the validation protocols: the validation against independent ground truth, the accuracy metrics calculation, the tolerance comparison, and the validation record. Three pressures bear on it. Time pressure arises when validation takes time and the consumer needs the product. Ground truth pressure arises when the ground truth is not available for the current run. Confidence pressure arises when the specialist is confident in the model and resents the validation. The specialist's validation discipline is the conduct of the validation on every run. The conduct is exercised at the validation steps in the flood model run and at the change detection validation (section 17.6).

The event discipline covers the quality standards under the time pressure of the event response. The event response is conducted under time pressure that the scheduled rhythm does not experience. Three pressures bear on it. Speed pressure arises when the emergency management authority needs the intelligence urgently and the specialist is tempted to waive the validation. Parameter pressure arises when the standard parameters may not be appropriate for the event's severity and the specialist is tempted to deviate without authorization. Documentation pressure arises when the documentation takes time and the specialist is tempted to defer it. The specialist's event discipline is the maintenance of the validation, the parameter authorization, and the documentation. The conduct is exercised at the flood event response (section 17.5), which is the gate through which the event-driven intelligence is produced.

The classification discipline covers the correct classification marking on every product, with particular attention to the Security Confidential marking for the energy infrastructure domain. Three pressures bear on it. Handling burden pressure arises when the Security Confidential marking requires more handling, which slows the release. Consumer access pressure arises when the consumer needs the product but does not have the clearance for the Security Confidential marking. Default pressure arises when the specialist is tempted to apply the lower marking to expedite the release. The specialist's classification discipline is the application of the correct classification marking to every product. The conduct is exercised at the classification marking verification (section 16.4), which the production operator conducts and the specialist supports.

A concrete example illustrates the disciplines' interaction. A flood event is developing in a monitored basin. The office manager activates the incident response, and the specialist receives the activation briefing: produce the flood inundation map, the exposure report, and the evacuation

guidance for the affected districts, with the Internal classification marking, within two hours. The specialist retrieves the latest sensor feed data and the latest processed imagery, configures the hydrology plug-in with the event-specific parameters (the observed rainfall, the observed river stages, the forecast rainfall), executes the model run, and validates the run against the observed event data. The specialist generates the products with the Internal classification marking, the model run's timestamp, and the data's cutoff time, and delivers the products to the office manager for review and release. The four disciplines are exercised in a single response: the modeling discipline (the event-specific parameters were documented), the validation discipline (the run was validated against the observed event), the event discipline (the validation was maintained under time pressure), and the classification discipline (the Internal marking was applied correctly).

The four disciplines are exercised through the specialist's daily rhythm and the office's procedures. The emergency management authority receives flood intelligence that is accurate, timely, and correctly classified. The energy sector entity receives infrastructure intelligence that is accurate, actionable, and security-confidential. The sectoral regulator receives the regulatory compliance products on the monthly cycle. The national audit authority examines the validation records and the classification compliance during the annual audit. The disciplines are established by the role family's training pipeline and qualifications framework, treated in Chapter 4, and maintained by the daily practice of the specialist's work.

17.8 Worked Example: A Flood Event Response with Data Quality Issues

At 03:00 on a Wednesday morning, the office manager activates the incident response for a flood event on the Riverbourne. The sensor feed dashboard shows the river stage at Gauge RB-04 at 3.8 metres (above the 5-year threshold of 3.5 metres) and rising, with 120 millimetres of rainfall in the preceding 6 hours. The manager's briefing to the specialist specifies: produce a flood inundation map, an exposure report, and evacuation guidance for the affected districts, with Internal classification, within 4 hours.

The specialist retrieves the latest sensor feed data from the processing engine. The rain gauge data from 14 of 15 gauges is current (within 15 minutes), but Gauge RG-11 is offline (the data pipeline operator has already reported this in the morning ingestion check). The river stage data from all 8 gauges is current. The specialist retrieves the latest processed imagery: a Sentinel-1 SAR image acquired at 05:30 UTC (which can be used for validation of the model's inundation extent), and the LiDAR DEM (2-metre resolution, acquired 2019). The specialist notes that the DEM does not reflect the 2022 Millbrook embankment (the technical lead has already documented this discrepancy in a previous model run, section 14.8), and the specialist knows that the model will over-predict the inundation in the southeastern portion of the basin near the Millbrook confluence.

The specialist configures the hydrology plug-in with the event-specific parameters. The rainfall data uses the 14 available gauges, with the rainfall at the RG-11 location interpolated from the surrounding gauges using the inverse distance weighting method. The river stage data from all 8 gauges is used as the boundary conditions. The HEC-HMS model uses the SCS curve number method with the event-specific antecedent moisture condition set to III (wet, based on the 180 millimetres of rainfall in the preceding 5 days, which exceeds the threshold for AMC III). The HEC-RAS 2D model uses the standard mesh (5-metre grid, with breaklines along the levees and road embankments) and

the standard Manning's n values (0.035 for the main channel, 0.060 for urban, 0.080 for forest, 0.030 for agriculture). The specialist documents the parameter changes from the standard set: the AMC changed from II to III, the rainfall interpolation for RG-11, and the known DEM discrepancy at the Millbrook embankment.

The specialist executes the model run. The HEC-HMS run completes in 3 minutes, producing a peak flow of 450 cubic metres per second at the catchment outlet (6 per cent higher than the standard-parameters run, due to the higher AMC). The HEC-RAS 2D run completes in 42 minutes, producing the inundation extent and depth at 5-metre grid resolution.

The specialist validates the model run against the observed data. The peak stage at Gauge RB-04 is predicted at 3.72 metres (observed 3.85 metres, difference 0.13 metres, within the 0.30-metre tolerance). The peak stage at Gauge RB-06 is predicted at 2.91 metres (observed 2.78 metres, difference 0.13 metres, within tolerance). The peak stage at Gauge RB-08 is predicted at 1.94 metres (observed 1.88 metres, difference 0.06 metres, within tolerance). The timing of the peak at all three gauges is within 30 minutes of the observed timing (within the 2-hour tolerance). The specialist compares the model's inundation extent with the Sentinel-1 SAR-derived flood extent: the true positive rate is 84 per cent in the non-affected areas (above the 80 per cent tolerance) and 72 per cent in the southeastern portion near the Millbrook embankment (below tolerance, but the discrepancy is known and documented). The overall true positive rate is 82 per cent, within tolerance.

The specialist generates the products. The flood inundation map shows the flood extent and depth, with the southeastern over-prediction noted in the metadata caveat ("The inundation extent in the southeastern portion of the basin, near the Millbrook confluence, may be over-predicted due to a flood defence embankment constructed in 2022 that is not reflected in the 2019 LiDAR DEM"). The exposure report identifies 347 affected parcels (23 residential, 4 commercial, 320 agricultural), 1,200 people within the inundated area, and 2 substations and 1 transmission corridor affected. The evacuation guidance identifies the evacuation routes and the shelter locations, with the caveat that the southeastern portion's evacuation guidance should be verified by the authority's field teams due to the known over-prediction. All products carry the Internal classification marking.

The specialist delivers the products to the office manager at 10:30 (within the 4-hour timeline), with the provenance record (the model run, the parameter changes, the validation results, the metadata caveat) and a verbal briefing on the data quality issues (the RG-11 gauge outage, the DEM discrepancy). The manager reviews the products, notes the metadata caveat, and authorizes the release at 11:15. The products are delivered to the emergency management authority through the secure data interface.

This worked example illustrates the hazard specialist's four disciplines operating together under event-driven pressure with data quality issues. The modeling discipline was exercised: the event-specific parameters were configured, documented, and justified. The validation discipline was exercised: the model was validated against 3 gauges and the SAR-derived extent, with the known discrepancy documented. The event discipline was exercised: the validation was maintained under the 4-hour deadline, and the metadata caveat was included to inform the consumer of the limitation. The classification discipline was exercised: the Internal marking was applied correctly. The specialist who maintains these disciplines under event-driven pressure with data quality issues is the specialist

whose intelligence is trusted by the emergency management authority, even when the intelligence includes caveats.

Chapter 18 — The Resource Economics Specialist

18.1 The Role of the Resource Economics Specialist

The resource economics specialist is the geospatial office's specialist for the food security and minerals domains. The specialist is accountable for the weekly crop bulletin, the strategic reserves procurement advisory, the minerals prospectivity, and the commercial commission delivery. The specialist's domain track spans the public-service and commercial sides of the office, and the specialist carries the commercial exposure that the firewall is designed to govern. The specialist operates the crop model and prospectivity plug-ins of the modeling workbench, treated in Chapter 10. The specialist reports to the technical lead, who configures the modeling workbench and validates the specialist's model runs, and performance is assessed by the office manager on the quality and timeliness of the specialist's products and on the specialist's adherence to the firewall.

The specialist's work produces the intelligence that the strategic reserve authority, the commercial minerals customers, and the strategic stockpile authority consume. The food security domain produces the weekly crop condition bulletin, the yield forecast, and the strategic reserves procurement advisory that the strategic reserve authority uses for the procurement planning. The minerals domain produces the prospectivity maps, the resource estimates, the critical minerals inventory, and the commercial minerals commissions that the commercial customers and the strategic stockpile authority use for the exploration planning and the strategic stockpile management.

The food security domain operates on a weekly cadence with seasonal milestones. The specialist produces the weekly crop condition bulletin every week, delivered by close of business each Wednesday, and updates the yield forecast as the season progresses. The consumer is the strategic reserve authority, and the classification regime is split: the crop condition summary is Public, and the detailed assessment and the yield forecast are Internal because their premature disclosure could move market prices.

The minerals domain operates on a project-based cadence with annual standing products. The specialist conducts the prospectivity model runs on a project basis for the commercial customers and annually for the standing products (the critical minerals inventory and the strategic stockpile assessment). The consumers are the commercial customers and the strategic stockpile authority, and the classification regime is mixed: Public for the base geology, Confidential for the commercial commissions, Security Confidential for the critical minerals inventory.

The specialist is the only officer in the office, apart from the office manager, who handles commercial products. The specialist produces the commercial minerals commissions under the commercial charge code, prepares the commercial release package, and presents the package to the office manager for the commercial release authorization (section 13.4). The specialist does not release the commercial product to the customer; the office manager authorizes the release, and the production operator releases the product through the commercial channel. The specialist's adherence to the firewall is verified by the office manager in the commercial release authorization and by the national audit authority in the annual audit.

The specialist's position is established at the operational tier of the role family, treated in Chapter 4. The position is held by an officer who has completed the training pipeline, holds the specialist certification in the resource economics domain, and has served a probationary period under the

supervision of a senior specialist. The position is a career position, and the specialist may progress to the senior specialist rank and then to the technical lead position with the additional certifications and experience specified by the qualifications framework. The specialist's training pipeline includes the foundational course (covering the pipeline architecture, the information system, the classification framework, and the office's procedures), the modeling workbench course (covering the workbench's data interface, parameter management, validation framework, and provenance recording), the crop model plug-in course (covering the vegetation indices, the crop growth models, the yield forecasting, and the validation against harvest outcomes), the prospectivity plug-in course (covering the weights-of-evidence and fuzzy logic methods, the machine learning classifiers, the resource estimation, and the validation against known deposits), and the commercial release procedure course (covering the firewall's four mechanisms, the commercial release log, the charge code management, and the classification marking for commercial products). The specialist must pass the certification assessment for each course before operating the corresponding plug-in or conducting the commercial release procedure. The specialist also receives training on the international reporting standards for mineral resources (the JORC Code, NI 43-101, or the SME Standard), because the minerals domain's resource estimates must comply with these standards when they are delivered to commercial customers who may use them in regulatory filings or investment decisions.

National mapping agencies organize the resource economics specialist function in comparable ways. The United States Geological Survey's agriculture specialists and economic geologists conduct the crop assessment and the minerals prospectivity for the Survey's programs, in coordination with the Department of Agriculture and the Department of the Interior. The European Commission's Joint Research Centre conducts the Monitoring of Agricultural Resources for Europe (MARS) and the minerals intelligence network, with agriculture specialists and economic geologists conducting the assessment and the prospectivity. Geoscience Australia's resource economics specialists conduct the minerals prospectivity and the critical minerals assessment for the agency's programs, in coordination with the Department of Industry, Science and Resources. The geospatial office's resource economics specialist conducts the crop assessment and the minerals prospectivity under the supervision of the technical lead, with the assessment and the prospectivity conducted to the technical lead's configuration and the model runs validated against independent ground truth before release.

18.2 The Specialist's Daily Rhythm

The resource economics specialist's day is organized around the variable modeling cadence of the two domains. The rhythm is anchored by the weekly crop condition bulletin (section 18.3), produced every week and delivered by close of business each Wednesday. The bulletin comprises the crop condition assessment, the yield forecast update, and a summary of the season's progress. The bulletin is conducted at the modeling workbench using the crop model plug-in, and it consumes the optical imagery, the weather data, the soil moisture data, the cadastre parcels, and the farmer-reported planting data.

The minerals prospectivity work (section 18.5) is conducted on a project basis, with the timelines agreed per project. The work comprises the retrieval of the geological, geophysical, and geochemical data for the project area; the configuration of the prospectivity plug-in; the execution of the model run; the validation against known deposits; and the generation of the prospectivity map, the

resource estimate, and (for standing products) the critical minerals inventory. The work takes from two weeks to six months per project, depending on the project's scope.

The strategic reserves procurement advisory (section 18.4) is produced as the yield forecast and the reserve stock levels warrant. The advisory synthesizes the latest yield forecast with the current reserve stock levels to produce a release or buy signal. The advisory is Internal because its premature disclosure could move market prices, and it is delivered to the strategic reserve authority through the internal distribution channel.

The commercial commission production (section 18.6) is conducted per the customer's commission. The specialist receives the commission from the office manager, conducts the prospectivity model run under the commercial charge code, produces the commercial product with the Confidential classification marking, prepares the commercial release package, and presents the package to the office manager for the commercial release authorization. The specialist does not release the product to the customer.

The specialist's rhythm is distinct from the hazard specialist's rhythm. The hazard specialist operates entirely on the public-service side; the resource economics specialist operates on both sides. The separation between the public-service and commercial work is maintained through the charge codes (separate codes for commercial and public-service work), the classification markings (Confidential for commercial products, Public or Internal for public-service products), and the commercial release procedure (the office manager's authorization for every commercial product). The separation is the operational expression of the firewall, treated in Chapter 5.

The daily rhythm is the specialist's mechanism for maintaining the analytical integrity of the two domains and the firewall's integrity in the commercial arm. A specialist who does not produce the weekly crop condition bulletin on time fails the strategic reserve authority. A specialist who does not maintain the separation between the commercial and public-service work breaches the firewall. A specialist who does not validate every model run produces intelligence that may be wrong in ways that are not visible to the consumer. The discipline is the daily practice of being present at the modeling workbench and informed about the two domains' states and the firewall's maintenance.

18.3 The Weekly Crop Condition Bulletin

The weekly crop condition bulletin is the resource economics specialist's recurring task for the food security domain, produced every week and delivered by close of business each Wednesday. The bulletin comprises the crop condition assessment, the yield forecast update, and a summary of the season's progress. The bulletin is conducted at the modeling workbench using the crop model plug-in, and it consumes the optical imagery, the weather data, the soil moisture data, the cadastre parcels, and the farmer-reported planting data. The bulletin is the office's weekly intelligence product for the strategic reserve authority.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the latest optical imagery for the agricultural regions.	Verify that the imagery was acquired within the bulletin's weekly window, and that the cloud masking has been applied to remove cloud-contaminated pixels.	The crop condition assessment depends on cloud-free imagery, and the weekly window ensures that the assessment reflects the current state of the crops.
2. Configure the crop model plug-in for the bulletin run.	Set the crop type, the growth stage, the weather data integration, and the soil moisture integration per the standard parameter set for the season and the region.	The crop model parameters determine the assessment's accuracy, and the parameters must be appropriate for the crop type and the growth stage, which change through the season.
3. Execute the crop condition assessment.	Calculate the vegetation indices (NDVI, NDWI) from the imagery, and stratify the indices by parcel using the cadastre feed to produce parcel-level condition assessments.	The parcel-level stratification is the mechanism that connects the crop condition to the specific fields the strategic reserve authority manages, and the connection is the condition of the assessment's usefulness.
4. Update the yield forecast.	Integrate the current condition assessment with the historical yield data and the seasonal progress to produce the updated yield forecast, and document the forecast's assumptions and uncertainties.	The yield forecast is the forward-looking intelligence that the strategic reserve authority uses for procurement planning, and the assumptions and uncertainties are the conditions of the forecast's interpretability.

Important Steps	Key Points	Reasons for Key Points
5. Validate the bulletin against ground truth.	Compare the condition assessment with the ground truth observations from the agricultural extension service, and calculate the accuracy metrics per the domain's validation protocol.	The validation is the analytical integrity guarantee of the food security domain, and the validation against ground truth is the check that confirms the assessment's accuracy.
6. Generate the bulletin with the Public classification marking for the summary and the Internal marking for the detailed assessment.	Apply the split classification regime: the crop condition summary is Public, and the detailed assessment and the yield forecast are Internal because their premature disclosure could move market prices.	The split classification regime is the security and economic integrity guarantee of the food security domain, and the correct marking is the condition of the bulletin's compliance.
7. Submit the bulletin to the production operator and record the run in the catalog and archive submodule.	Submit the bulletin with the provenance record and the validation results, and record the run with the model version, parameters, and validation results.	The submission and the recording complete the weekly cycle, and the provenance and the validation are the integrity guarantees that the production operator and the archive rely on.

18.4 The Strategic Reserves Procurement Advisory

The strategic reserves procurement advisory is the resource economics specialist's task of producing the advisory that synthesizes the yield forecast with the current reserve stock levels to produce a release or buy signal. The advisory is the office's forward-sensor product for the strategic reserve authority. The advisory is Internal because its premature disclosure could move market prices, and the advisory is delivered to the reserve authority through the internal distribution channel.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the latest yield forecast from the crop model.	Verify that the yield forecast is the current forecast, and that the forecast's assumptions and uncertainties are documented.	The procurement advisory synthesizes the yield forecast, and the advisory's reliability depends on the forecast's currency and documentation.

Important Steps	Key Points	Reasons for Key Points
2. Retrieve the current reserve stock levels from the strategic reserve authority.	Obtain the stock levels through the data-sharing arrangement with the reserve authority, and verify the stock levels' date and source.	The procurement advisory synthesizes the yield forecast with the stock levels, and the stock levels must be current for the synthesis to be meaningful.
3. Analyze the supply-demand balance.	Compare the forecast supply (the yield forecast) with the expected demand (the historical consumption adjusted for population and market trends) and the current stock levels, and identify whether the balance indicates a surplus (release signal) or a deficit (buy signal).	The supply-demand analysis is the analytical core of the advisory, and the analysis must be grounded in the forecast and the stock levels, not in the specialist's prior expectations.
4. Produce the release or buy signal with the supporting analysis.	State the signal clearly (release, buy, or hold), and provide the supporting analysis that justifies the signal, including the forecast, the stock levels, the demand assumptions, and the uncertainties.	The signal is the advisory's conclusion, and the supporting analysis is the advisory's justification; the reserve authority decides, and the office's role is to provide the analysis that informs the decision, not to make the decision.
5. Apply the Internal classification marking to the advisory.	Mark the advisory as Internal, and include the handling caveat that restricts the advisory to the reserve authority and the office.	The Internal marking is the economic integrity guarantee, and the handling caveat is the mechanism that prevents the advisory's premature disclosure, which could move market prices.

Important Steps	Key Points	Reasons for Key Points
6. Submit the advisory to the production operator for production and delivery.	Submit the advisory with the provenance record, and specify the internal distribution channel for delivery to the reserve authority.	The internal distribution channel is the mechanism that enforces the Internal marking, and the provenance record is the integrity guarantee that the advisory is traceable to its inputs.
7. Record the advisory in the catalog and archive submodule.	Record the advisory's identifier, the yield forecast version, the stock level date, the signal, and the Internal classification marking.	The advisory record is the audit trail for the food security domain's forward-sensor function, and it is the record that the national audit authority examines during the annual audit.

18.5 The Minerals Prospectivity Model Run

The minerals prospectivity model run is the resource economics specialist's task for the minerals domain, conducted on a project basis for commercial customers and annually for the standing products. The run produces the prospectivity maps, the resource estimates, and the critical minerals inventory that the minerals domain delivers. The run is conducted at the modeling workbench using the prospectivity plug-in, and it consumes the geological survey data, the aeromagnetic and gravity geophysics, the geochemistry, the hyperspectral imagery, and the historical exploration data. The run is conducted under the technical lead's configuration and validated against known deposits.

Important Steps	Key Points	Reasons for Key Points
1. Retrieve the geological, geophysical, and geochemical data for the project area.	Verify that the data's provenance record is complete, and that the data's license terms permit the intended use (commercial or public-service).	The prospectivity model's accuracy depends on the completeness of the input data, and the license terms determine whether the data may be used for the commercial product.

Important Steps	Key Points	Reasons for Key Points
2. Configure the prospectivity plug-in for the model run.	Set the prospectivity model type (weights-of-evidence or machine learning), the training data, the predictor variables, and the validation set per the standard parameter set for the mineral type and the geological setting.	The prospectivity model parameters determine the model's behavior, and the parameters must be appropriate for the mineral type and the geological setting, which vary by project.
3. Execute the prospectivity model run.	Monitor the model's convergence and the predictor variable contributions, and note any anomalies that may indicate a data quality issue or a model overfit.	The model's convergence and the variable contributions are the indicators of the model's reliability, and the anomalies are the early warning signs that the specialist must investigate.
4. Validate the model run against known deposits.	Compare the model's prospectivity predictions with the locations of known deposits in the project area, and calculate the accuracy metrics per the domain's validation protocol.	The validation against known deposits is the analytical integrity guarantee of the minerals domain, and the model run that fails validation must not be released.
5. Generate the prospectivity map, the resource estimate, and (for standing products) the critical minerals inventory.	Generate the products with the classification marking determined by the product type: Public for base geology, Confidential for commercial commissions, Security Confidential for the critical minerals inventory.	The mixed classification regime is the hardest records-management problem in the office, and the correct marking is the condition of the products' compliance and security.

Important Steps	Key Points	Reasons for Key Points
6. For commercial commissions, apply the commercial charge code and prepare the product for the commercial release procedure.	Verify that the product was produced from commercial data sources and commercial model runs, and that the commercial charge code is applied; do not release the product directly to the customer.	The firewall requires commercial products to be produced from commercial sources and released through the commercial release procedure, and the specialist's preparation is the first step in the procedure that the office manager completes.
7. Submit the product to the production operator and record the run in the catalog and archive submodule.	Submit the product with the provenance record and the validation results, and record the run with the model version, parameters, validation results, and the charge code (commercial or public-service).	The submission and the recording complete the model run, and the charge code is the financial mechanism of the firewall that the audit trail must reflect.

18.6 The Commercial Commission Production and Release

The commercial commission production and release is the resource economics specialist's task of producing and releasing a commercial minerals product to a commercial customer. The task spans the specialist's model run and the office manager's commercial release authorization, and it is the firewall in operation from the specialist's perspective. The specialist produces the product under the commercial charge code, prepares the product for the commercial release procedure, and presents the product to the office manager for authorization. The specialist does not release the product to the customer; the office manager authorizes the release, and the production operator releases the product through the commercial channel.

Important Steps	Key Points	Reasons for Key Points
1. Receive the commercial commission from the office manager.	Confirm the customer identity, the commission scope, the agreed price, the delivery timeline, and the commercial charge code; do not begin the work without the office manager's authorization.	The office manager is the sole authority for commercial terms, and the specialist who begins work without the manager's authorization has breached the firewall's authority mechanism.

Important Steps	Key Points	Reasons for Key Points
2. Conduct the prospectivity model run under the commercial charge code.	Use only commercial data sources and commercial model parameters; do not use public-service data or public-service model runs for the commercial commission.	The firewall requires commercial products to be produced from commercial sources, and the use of public-service data or models is the firewall's primary failure mode.
3. Produce the commercial product with the Confidential classification marking.	Apply the Confidential marking with the customer identity and the product identifier, and include the handling caveat that restricts the product to the customer and the office.	The Confidential marking is the classification mechanism of the firewall, and the customer identity is the restriction that prevents disclosure to non-minerals staff.
4. Prepare the commercial release package for the office manager.	Assemble the product, the provenance record (showing the commercial data sources and the commercial charge code), the validation results, the customer identity, the agreed price, and the commercial charge code confirmation.	The commercial release package is the information the office manager needs to conduct the commercial release authorization (Chapter 13, section 13.4), and the package must be complete for the manager to verify the firewall compliance.
5. Present the commercial release package to the office manager.	Brief the manager on the product, the firewall compliance, the classification marking, the customer identity, and the price; answer any questions the manager has; do not release the product to the customer under any circumstances.	The office manager is the sole authority for commercial releases, and the specialist who releases the product directly to the customer has breached the firewall's authority mechanism and undermined the firewall's integrity.

Important Steps	Key Points	Reasons for Key Points
6. After the manager's authorization, hand the product to the production operator for release.	Confirm that the manager has signed the commercial release log, and direct the production operator to release the product through the commercial channel.	The manager's signature on the commercial release log is the authorization, and the production operator's release through the commercial channel is the distribution mechanism of the firewall.
7. Record the commercial commission in the catalog and archive submodule.	Record the product identifier, the customer identity, the price, the commercial charge code, the model run provenance, the validation results, and the commercial release log entry.	The commercial commission record is the audit trail for the commercial arm, and it is the record that the national audit authority examines during the annual audit of the firewall's maintenance.

18.7 The Specialist's Discipline

The resource economics specialist's discipline is the daily practice of maintaining the analytical integrity of the food security and minerals domains and the firewall's integrity in the commercial arm, under operational pressure. The discipline has four components: modeling, validation, classification, and firewall. Each component is maintained through a specific procedure that surfaces the pressure and verifies the standard.

The modeling discipline covers the crop model and prospectivity model parameter sets, the model version control, the model run documentation, and the model provenance record, which are configured by the technical lead and applied by the specialist. Three pressures bear on it. Parameter pressure arises when the standard parameters are wrong for the season or the geological setting and the specialist is tempted to deviate without authorization. Version pressure arises when the latest model version is untested and the specialist is tempted to use it without the validation. Documentation pressure arises when the documentation takes time and the specialist is tempted to omit it. The specialist's modeling discipline is the application of the standard parameters, the use of the authorized model versions, and the documentation of every run. The conduct is exercised at the weekly crop condition bulletin (section 18.3) and the minerals prospectivity model run (section 18.5).

The validation discipline covers the validation protocols: the validation against independent ground truth (the agricultural extension service's observations for the crop model, the known deposits for the prospectivity model), the accuracy metrics calculation, the tolerance comparison, and the validation record. Three pressures bear on it. Time pressure arises when validation takes time and the consumer needs the product. Ground truth pressure arises when the ground truth is not available for the current run. Confidence pressure arises when the specialist is confident in the model and resents the validation. The specialist's validation discipline is the conduct of the validation on every run. The

conduct is exercised at the validation steps in the weekly crop condition bulletin and the minerals prospectivity model run.

The classification discipline covers the mixed classification regime: the Public marking for the crop condition summary and the base geology, the Internal marking for the detailed assessment and the yield forecast and the strategic reserves procurement advisory, the Confidential marking for the commercial minerals commissions, and the Security Confidential marking for the critical minerals inventory. Three pressures bear on it. Handling burden pressure arises when the higher markings require more handling, which slows the release. Consumer access pressure arises when the consumer needs the product but does not have the clearance for the higher marking. Default pressure arises when the specialist is tempted to apply the lower marking to expedite the release. The specialist's classification discipline is the application of the correct classification marking to every product. The conduct is exercised at the classification marking verification (section 16.4), which the production operator conducts and the specialist supports.

The firewall discipline covers the separation between the public-service and commercial work, maintained through the charge codes (separate codes for commercial and public-service work), the data sources (commercial data for commercial products, public-service data for public-service products), the model runs (separate runs for commercial and public-service products), and the release procedure (the office manager's authorization for every commercial product). Three pressures bear on it. Commercial pressure arises when a commercial customer needs the product urgently and the public-service data is available. Workload pressure arises when the specialist is overloaded and the commingling of the work is expedient. Financial pressure arises when the commercial arm's surplus is below target and the commingling inflates the surplus. The specialist's firewall discipline is the refusal to commingle the commercial and public-service work. The conduct is exercised at the commercial commission production and release (section 18.6), which is the gate through which the commercial product is verified for firewall compliance.

A concrete example illustrates the disciplines' interaction. A commercial customer commissions a prospectivity study for a critical minerals project, with a delivery deadline of six weeks and an agreed price. The specialist receives the commission from the office manager, confirms the commercial charge code, and conducts the prospectivity model run using only commercial data sources (the commercial geological survey, the commercial aeromagnetic data, the commercial hyperspectral imagery). The specialist validates the run against the known deposits in the project area, generates the prospectivity map and the resource estimate with the Confidential classification marking (including the customer identity and the product identifier), and prepares the commercial release package. The specialist presents the package to the office manager, who verifies the firewall compliance (commercial data sources, commercial charge code), the classification marking (Confidential with customer identity), the customer identity, and the agreed price, and signs the commercial release log. The four disciplines are exercised in a single commission: the modeling discipline (the standard parameters were applied and documented), the validation discipline (the run was validated against known deposits), the classification discipline (the Confidential marking was applied with the customer identity), and the firewall discipline (the commercial data sources and charge code were verified).

The four disciplines are exercised through the specialist's daily rhythm and the office's procedures. The strategic reserve authority receives the weekly crop condition bulletin on time, the procurement advisory when warranted, and the intelligence with the correct Internal marking. The commercial customers receive the prospectivity products that are accurate, confidential, and produced from commercial sources. The strategic stockpile authority receives the critical minerals inventory with the Security Confidential marking. The national audit authority examines the firewall's maintenance, the validation records, and the classification compliance during the annual audit. The disciplines are established by the role family's training pipeline and qualifications framework, treated in Chapter 4, and maintained by the daily practice of the specialist's work.

18.8 Worked Example: A Week in the Life of the Resource Economics Specialist

The resource economics specialist's week illustrates the dual-domain, dual-cadence, dual-classification character of the role. The week spans the public-service food security domain (weekly cadence, split classification: Public summary and Internal advisory) and the commercial minerals domain (project-based cadence, Confidential commercial commission), and the specialist must maintain the firewall between the two throughout.

Monday. The specialist begins the week by reviewing the crop condition data for the agricultural region. The data pipeline operator has processed the Sentinel-2 image acquired on Saturday (surface reflectance, cloud masking at 60 per cent confidence threshold, 15 per cent cloud cover). The specialist retrieves the processed data from the modeling workbench, configures the crop model plug-in with the standard parameters for the current growth stage (the crops are in the stem elongation stage, with the accumulated growing degree days at 850, indicating the transition from tillering to stem elongation), and computes the NDVI and EVI for each parcel using the cadastre feed for stratification. The specialist compares the current NDVI with the historical distribution for the same parcel, the same crop, and the same phenological stage: the condition score is at the 62nd percentile (above-normal condition, consistent with the favourable rainfall in the preceding two weeks). The specialist updates the yield forecast by integrating the condition score with the weather data (the seasonal rainfall is 15 per cent above normal, the temperature is within the optimal range for the crop's growth stage) and the historical yield data (the 10-year average yield for the region is 4.2 tonnes per hectare, and the current forecast is 4.5 tonnes per hectare, 7 per cent above the average). The specialist documents the model run in the provenance record.

Tuesday. The specialist completes the crop condition bulletin. The bulletin comprises the crop condition map (the green-to-red condition gradient, with the parcel boundaries overlaid), the yield forecast table (by district and by region), and the seasonal progress summary. The specialist applies the split classification marking: the crop condition summary is Public (for public awareness), and the detailed assessment and the yield forecast are Internal (because the premature disclosure of the yield forecast could move market prices). The specialist submits the bulletin to the production operator for production.

Wednesday. The bulletin is produced, reviewed by the office manager, and delivered to the strategic reserve authority by close of business. The specialist then turns to the commercial minerals commission. Meridian Resources has commissioned a prospectivity assessment for the northern

highlands (commission COMM-2024-017, $45,000, 8-week delivery). The specialist is in week 5 of the commission. The specialist has already completed the data integration (the commercial geological data, the commercial aeromagnetic data, and the public geological map have been ingested and processed) and the weights-of-evidence model run (the training data from the geological survey's mineral occurrence database has been used to compute the W+ and W- weights for each geological feature). The specialist is now conducting the validation: the model's prospectivity map is compared with the known deposits in the study area, and the accuracy metrics are calculated. The prospectivity accuracy is 78 per cent (above the 70 per cent tolerance), and the prediction ratio is 4.2 (above the 3.0 tolerance). The specialist documents the validation results in the provenance record, under the commercial charge code COMM-2024-017.

Thursday. The specialist produces the prospectivity map and the commission report. The prospectivity map uses the warm-cool gradient (dark red for the highest prospectivity, dark blue for the lowest), with the geological formations overlaid in muted browns and the known deposits marked as black stars. The commission report documents the methodology (the weights-of-evidence method, the training data, the W+ and W- weights for each feature), the validation results (78 per cent accuracy, 4.2 prediction ratio), and the resource estimate (the estimated quantity and quality of the mineral resource, following the JORC Code's inferred resource category). The specialist applies the Confidential classification marking, with the customer identity ("Meridian Resources") and the product identifier ("COMM-2024-017") in the marking. The specialist submits the commission report to the production operator for production.

Friday. The commission report is produced, and the specialist prepares the commercial release package for the office manager. The package includes the product (the prospectivity map and the commission report), the provenance record (showing the commercial data sources and the commercial charge code), the validation results, the customer identity (verified against the commercial customer register), and the agreed price ($45,000, verified against the pricing schedule). The office manager reviews the package under the commercial release authorization procedure (section 13.4): the firewall compliance is verified (the product was produced from commercial data sources and commercial model runs, charge code COMM-2024-017), the classification marking is verified (Confidential with the customer identity), and the commercial release log is signed. The production operator releases the product to Meridian Resources through the commercial distribution channel.

The specialist's week illustrates the firewall's operation in the daily work of the officer who carries the commercial exposure. The specialist produced a public-service product (the crop condition bulletin, Public/Internal, public-service charge code PUB-2024-120) and a commercial product (the prospectivity map and commission report, Confidential, commercial charge code COMM-2024-017) in the same week, using different data sources (public Sentinel-2 imagery vs. commercial geological data), different model runs (the crop model plug-in vs. the prospectivity plug-in), different charge codes, different classification markings, and different distribution channels. The two products did not commingle at any stage of the pipeline, and the firewall was maintained throughout. The specialist who maintains this separation in the daily work is the specialist whose public-service intelligence is trusted by the strategic reserve authority and whose commercial

intelligence is trusted by the commercial customers, and the trust on both sides is the condition of the revenue-neutral model's sustainability.

Chapter 19 — Flood Risk Intelligence

Part IV — The Four Domains: The four geospatial domain applications, each treated as a doctrine chapter with definition, scope, method variations, products, consumers, cadence, and classification regime.

19.1 The Domain Defined

Flood risk intelligence is the production of maps, models, and reports describing the probability, extent, depth, and consequences of flooding across a jurisdiction's territory. Flooding is the most common natural hazard worldwide, affecting an estimated 1.8 billion people between 2000 and 2019 and causing approximately $651 billion in economic losses over the same period, according to the Centre for Research on the Epidemiology of Disasters. The domain serves the emergency management authority as its principal consumer. Its products — flood inundation maps, exposure reports, and evacuation guidance — support flood preparedness, flood response, and flood recovery. The domain operates on a continuous cadence with event-driven production, under a tiered classification regime that ranges from Public for general flood hazard maps to Security Confidential for scenarios involving dam or levee failure.

Flooding takes several forms, and the domain must address each. Riverine flooding occurs when the flow in a river exceeds the channel's capacity and the water spreads onto the floodplain: the 2021 European floods (July 2021, Germany, Belgium, Netherlands) caused approximately $46 billion in damages and 240 deaths, and demonstrated the catastrophic potential of riverine flooding when extreme rainfall coincides with saturated soils. Coastal flooding occurs when storm surge, high tides, or sea-level rise inundate low-lying coastal areas: Hurricane Katrina (August 2005, United States) caused approximately $125 billion in damages and over 1,800 deaths, and demonstrated the catastrophic potential of coastal flooding when storm surge overtops levees. Flash flooding occurs when intense rainfall over a short period overwhelms the drainage capacity of a small catchment: the 2021 Ahr valley flood (July 2021, Germany) caused 196 deaths and demonstrated the speed at which flash flooding can develop and the inadequacy of riverine flood models for small, steep catchments. Dam and levee failure flooding occurs when a hydraulic structure fails, releasing a large volume of water downstream: the Banqiao Dam failure (August 1975, China) caused an estimated 171,000 deaths and demonstrated the catastrophic potential of structural failure.

The domain is bounded by three interfaces. The flood-water boundary separates it from the water authority's water governance function (Chapter 27): the geospatial office models where water will go; the water authority manages the waterways that convey it. The intelligence-response boundary separates it from the emergency management authority's response function (Chapter 24): the geospatial office produces the intelligence; the authority conducts the response. The recording-analyzing boundary separates it from the cadastral authority (Chapter 23): the cadastral authority records the land; the geospatial office analyzes the flood risk to the land.

19.2 The Scope

The flood risk domain covers three activities: flood hazard assessment, flood exposure assessment, and flood event response.

Flood hazard assessment is the modeling of the probability, extent, and depth of flooding across the jurisdiction's territory. The probability is expressed as a return period: the 100-year flood (the flood with a 1 per cent annual exceedance probability) is the standard regulatory benchmark in many jurisdictions, but the domain also models the 10-year, 50-year, 500-year, and the probable maximum flood (PMF) for dam safety assessments. The extent is the geographic area inundated by the flood, and the depth is the vertical distance from the ground surface to the water surface. The hazard assessment is conducted through coupled hydrological and hydraulic models, treated in Chapter 10, section 10.2.

Flood exposure assessment is the analysis of the population, property, and infrastructure exposed to the flood hazard. The exposure assessment overlays the flood hazard maps with the cadastre (to identify the affected parcels), the population data (to estimate the affected population), and the infrastructure asset registry (to identify the affected infrastructure). The exposure assessment is the mechanism that makes the hazard assessment actionable: a flood hazard map without the exposure assessment is a map of where the water goes, but the emergency management authority needs to know who and what is in the water's path.

Flood event response is the production of real-time intelligence during a flood event. The event response is activated when the sensor feeds indicate that a flood is imminent or occurring, and it overrides the scheduled cadence for the duration of the event. The event response produces real-time inundation maps (based on the latest sensor feed data and the current model run), updated exposure reports (based on the real-time inundation), and evacuation guidance (based on the inundation extent and the road network).

The scope does not include the operational response to the flood event, which is the emergency management authority's function. It does not include waterway management (levee maintenance, channel dredging, gate operation), which is the water authority's function. It does not include land-use planning that determines floodplain development, which is the spatial planning authority's function. An office that expands its scope to include operational response, waterway management, or land-use planning exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

19.3 The Method

The method of the flood risk domain is coupled hydrological and hydraulic modeling. The hydrological model simulates the rainfall-runoff process, converting the rainfall data into a hydrograph (a time series of the river flow at the catchment outlet). The hydraulic model simulates the channel flow and the floodplain inundation, converting the hydrograph into an inundation extent and depth map. The two models are coupled: the hydrological model's output is the hydraulic model's input.

The hydrological model takes the rainfall data (from rain gauges, weather radar, or satellite precipitation estimates), the soil data (from the soil map, classified by the hydrologic soil group: A for high infiltration, B for moderate, C for slow, D for very slow), the land cover data (from the land cover base layer, classified by the SCS curve number), and the topographic data (from the DEM, for the catchment delineation and the flow direction), and produces the hydrograph. The rainfall-runoff transformation methods include the SCS curve number method (used for design flood estimation and

for many operational flood forecasting systems), the Green-Ampt infiltration model (used for physically-based distributed models), and the Sacramento Soil Moisture Accounting model (used by the US National Water Model for continuous simulation).

The hydraulic model takes the hydrograph as the upstream boundary condition, the river stage from the sensor feeds as the downstream boundary condition, the channel geometry from the hydrography base layer and the DEM, and the roughness coefficients (Manning's n values) from the land cover data, and produces the water surface elevation, the flow velocity, and the flow depth at each point in the channel and the floodplain. The one-dimensional models (HEC-RAS 1D, MIKE 11) are computationally efficient and suitable for long river reaches where the floodplain flow is predominantly parallel to the channel. The two-dimensional models (HEC-RAS 2D, MIKE 21, TUFLOW) are computationally intensive and suitable for urban areas and complex floodplains where the flow is multidirectional. The coupled 1D-2D models (HEC-RAS 1D-2D, MIKE 11-21) combine the efficiency of the 1D model for the channel with the accuracy of the 2D model for the floodplain. The coupled approach is particularly useful for large river basins where the channel is well-defined (and the 1D model is sufficient) but the floodplain includes urban areas where the flow is multidirectional (and the 2D model is necessary). The coupling is implemented through lateral weirs that connect the 1D channel to the 2D floodplain: the 1D model computes the water surface elevation in the channel, and when the elevation exceeds the bank elevation, the water spills laterally into the 2D floodplain, where the 2D model computes the flow's direction, depth, and velocity. The coupled model's computational cost is between the 1D and the 2D models, and the coupled model is the standard approach for the office's flood hazard assessments of large basins with urban floodplains. The model's accuracy depends on the input data's quality. The DEM's resolution and vertical accuracy are the most critical factors: a LiDAR DEM at 1-metre resolution with 0.10-metre vertical RMSE produces inundation maps that resolve individual levees and road embankments, while a SRTM DEM at 30-metre resolution with 5-metre vertical RMSE produces inundation maps that smooth over these features and over-predict the inundation extent in flat areas. The Manning's n values are the second most critical factor: the values are derived from the land cover, but the derivation involves judgment (a forested floodplain's roughness depends on the tree density, the undergrowth, and the flood depth), and the values are typically calibrated against observed events. The boundary conditions are the third most critical factor: the upstream boundary (the hydrograph from the hydrological model) and the downstream boundary (the river stage from the sensor feeds or the tidal gauge) must be accurate, because the model's results are sensitive to the boundary conditions, particularly in the tidal reach where the riverine flood and the storm surge interact.

The method includes the exposure assessment, which overlays the flood hazard maps with the cadastre (identifying the affected parcels by their parcel identifiers), the population data (estimating the affected population from the census data, distributed by the parcel or the building footprint), and the infrastructure asset registry (identifying the affected substations, transmission corridors, water treatment plants, hospitals, and emergency facilities). The exposure assessment is the mechanism that makes the flood hazard maps actionable for the emergency management authority: the authority's evacuation decisions are based not on the flood extent alone but on the population and the infrastructure within the flood extent.

The method also includes the validation against observed flood events. The validation compares the model's predicted inundation extent with the observed extent (from satellite imagery, particularly Sentinel-1 SAR, which can map the flood extent through clouds) and the model's predicted depth with the observed depth (from the gauge network's recorded stages and from high-water marks surveyed after the event). The validation is the analytical integrity guarantee of the flood risk domain, and a model that has not been validated against observed events produces inundation maps of unknown accuracy.

19.4 The Products

The flood risk domain produces four product types, each serving a different phase of the flood management cycle.

Flood hazard maps show the probability, extent, and depth of flooding for the standard return periods (typically the 10-year, 50-year, 100-year, and 500-year floods). The maps are the domain's standing product, updated seasonally or when new data (a new DEM, a recalibrated model) warrants a revision. The maps are used by the spatial planning authority for floodplain development planning, by the insurance industry for flood insurance pricing, and by the public for flood awareness. The maps carry the Public classification, because public awareness of the flood hazard is the condition of public preparedness.

Exposure reports quantify the population, property, and infrastructure exposed to the flood hazard. The reports are produced for the emergency management authority (for preparedness and recovery planning) and for the sectoral regulators (for the infrastructure's resilience assessment). The reports include the parcel-level exposure (which parcels are within the flood hazard area, and what is their land use, their value, and their occupancy), the population exposure (how many people live within the flood hazard area, disaggregated by age, income, and vulnerability), and the infrastructure exposure (which substations, transmission corridors, water treatment plants, hospitals, and emergency facilities are within the flood hazard area). The reports carry the Internal classification when they include critical infrastructure exposure, because the disclosure of the critical infrastructure's vulnerability could inform attacks.

Evacuation guidance identifies the evacuation routes and the shelter locations for the flood hazard. The guidance is produced for the emergency management authority (for the evacuation planning) and for the public (for the individual preparedness). The guidance includes the evacuation routes (the roads that remain passable during the flood, identified by overlaying the flood hazard map with the road network and the road elevation), the shelter locations (the buildings that are outside the flood hazard area and that have the capacity to accommodate the evacuated population), and the evacuation timing (the lead time required for the evacuation, based on the population, the evacuation routes' capacity, and the flood's arrival time). The guidance carries the Public classification, because the public's ability to evacuate depends on the public's knowledge of the routes and the shelters.

Event-specific intelligence is produced during a flood event. It comprises real-time inundation maps (based on the latest model run, with the current sensor feed data as the boundary conditions), event exposure reports (based on the real-time inundation, with the current population and infrastructure data), and event evacuation guidance (based on the real-time inundation, with the current road network status). The event-specific intelligence carries the Internal classification

(because it includes the critical infrastructure exposure) or the Public classification (for the evacuation guidance distributed to the public through the emergency management authority). The emergency management authority uses the event-specific intelligence for the operational response: the evacuation orders, the resource deployment, and the public communication.

19.5 The Consumers

The flood risk domain serves four consumers, each with different needs and different operational contexts.

The emergency management authority is the principal consumer. The authority uses the flood hazard maps for preparedness (the development of evacuation plans, the pre-positioning of resources, the public education), the event-specific intelligence for response (the evacuation orders, the resource deployment, the situational awareness), and the exposure reports for recovery (the damage assessment, the recovery planning, the mitigation). The authority requires the intelligence in operational formats (GeoTIFF for its GIS system, PDF for its briefing packs) and on event-driven timelines (within hours for the event-specific intelligence). The authority's relationship with the geospatial office is governed by the intelligence-response division (Chapter 24): the office provides the intelligence, and the authority conducts the response.

The water authority is a secondary consumer. The authority uses the flood hazard maps for the waterway management: the levee maintenance (the maps identify the levees that are overtopped by the design flood), the channel dredging (the maps identify the channels where the sediment accumulation has reduced the capacity), and the gate operation (the maps identify the gates that must be operated during the flood). The authority requires the model integration (the flood model's results must be compatible with the authority's own hydraulic models) and the data sharing (the sensor feed data must be shared in real time). The authority's relationship with the geospatial office is governed by the flood-water boundary (Chapter 27): the office models where water will go, and the authority manages the waterways that convey it.

The spatial planning authority is a tertiary consumer. The authority uses the flood hazard maps for the floodplain development planning: the zoning (the maps identify the areas where development is prohibited, restricted, or permitted with conditions), the building codes (the maps identify the flood-resistant construction requirements for each zone), and the development permitting (the maps are the basis for the development permit's flood risk assessment). The authority requires the hazard maps in a format that can be incorporated into the zoning ordinances and the building codes, typically as a flood hazard overlay district.

The public is a fourth consumer. The public uses the public flood hazard maps for the individual preparedness: the awareness of the flood risk for one's property, the purchase of flood insurance, the preparation of an evacuation plan, and the monitoring of the flood warnings. The public requires the maps in an accessible format (a web map, a printed brochure, or a mobile application) and in a language that the public understands.

19.6 The Cadence

The flood risk domain operates on a continuous cadence with event-driven production. The continuous cadence comprises the continuous ingestion of sensor feeds (rain gauges at 15-minute

intervals, river stage sensors at 15-minute intervals, weather stations at hourly intervals), the daily processing of the acquired satellite imagery (Sentinel-2 for the optical monitoring of the catchment conditions, Sentinel-1 for the radar monitoring of the soil moisture), the daily review of the flood risk by the hazard specialist (a brief assessment of the sensor feed data, the weather forecast, and the flood risk for the monitored basins), and the seasonal updating of the flood hazard maps (the recalibration of the hydrological model based on the events observed during the season, the updating of the hydraulic model based on any changes in the channel geometry or the land cover).

The event-driven production comprises the activation of the event response when a flood event is imminent or occurring (triggered by the sensor feed alerts or by a request from the emergency management authority), the production of the event-specific intelligence during the event (the real-time model run with the current sensor feed data, the exposure assessment, the evacuation guidance, and the delivery to the emergency management authority within the agreed timeline), the continuous updating of the intelligence as the event evolves (the model run is repeated with the latest sensor feed data, and the updated intelligence is delivered to the authority), and the stand-down of the response when the event is resolved (the river stages return to below the flood threshold, the rainfall has ceased, and the emergency management authority confirms that the response is no longer required).

The continuous cadence is maintained by the data pipeline operator, who ingests the sensor feeds and processes the satellite imagery, and the hazard specialist, who monitors the feeds and conducts the daily flood risk review. The event-driven production is maintained by the office manager, who activates the response and coordinates with the emergency management authority, and the hazard specialist, who produces the event-specific intelligence. During a flood event, the event-driven cadence overrides the scheduled cadence, and the office's work is organized around the incident response activation (section 13.6).

19.7 The Classification Regime

The flood risk domain operates under a tiered classification regime that reflects the varying sensitivity of the domain's products.

General flood hazard maps are Public. Public awareness of the flood hazard is the condition of public preparedness, and public preparedness is the condition of public safety. The United States' National Flood Insurance Program makes its Flood Insurance Rate Maps publicly available through the FEMA Flood Map Service Center, and the European Floods Directive requires member states to make their flood hazard maps publicly available. The Public classification ensures that the public can assess the flood risk for their property, purchase flood insurance, and prepare evacuation plans. The FEMA Flood Map Service Center provides the FIRMs through a web interface that allows any member of the public to search for their property's address and view the flood hazard map for their area, and the European Floods Directive requires the member states to provide their flood hazard maps through a publicly accessible web portal. The public availability of the flood hazard maps is a policy choice: the maps could be restricted to the government and the emergency management authority, but the policy choice is to make them public, because the public's awareness of the flood risk is the condition of the public's preparedness, and the public's preparedness reduces the flood's impact on life and property.

Exposure reports for critical infrastructure are Internal. The exposure of critical infrastructure — substations, transmission corridors, water treatment plants, hospitals, emergency facilities — is sensitive, because the disclosure of the infrastructure's vulnerability could inform attacks. The Internal classification restricts the exposure reports to the authorized government personnel (the emergency management authority, the sectoral regulators, the infrastructure operators), while ensuring that the personnel who need the information for the preparedness and the response have access to it.

Scenarios involving dam or levee failure are Security Confidential. The disclosure of the dam failure inundation maps could inform attacks on the dams, because the maps identify the dams whose failure would cause the most catastrophic flooding downstream. The Security Confidential classification restricts the dam failure scenarios to the security-cleared personnel (the dam operators, the emergency management authority's security-cleared staff, the intelligence agencies). The United States' Dam Safety Act of 2006 (Public Law 109-460) restricts the public release of dam failure inundation maps for high-hazard dams, and the classification is consistent with the Security Confidential level of the geospatial office's classification framework.

19.8 The Comparative Practice

The United States National Flood Insurance Program (NFIP), established by the National Flood Insurance Act of 1968 and administered by the Federal Emergency Management Agency (FEMA), is the world's largest flood hazard mapping programme. The NFIP produces Flood Insurance Rate Maps (FIRMs) that delineate the Special Flood Hazard Areas (the 100-year floodplain) for approximately 22,000 communities participating in the programme. The FIRMs are the basis for the federal flood insurance requirements (the mandatory purchase requirement for properties in the Special Flood Hazard Areas) and for the local floodplain development regulations (the minimum building standards for structures in the floodplain). The FIRMs are produced through FEMA's Risk Mapping, Assessment, and Planning (Risk MAP) programme, which uses FEMA's flood hazard models (typically HEC-RAS for the riverine flooding and ADCIRC for the coastal storm surge) and FEMA's mapping standards (the Guidelines and Standards for Flood Risk Analysis and Mapping). The FIRMs are publicly available through the FEMA Flood Map Service Center, and the programme's Public classification of the flood hazard maps is the exemplar of the public-awareness rationale for the Public classification level.

The European Floods Directive (Directive 2007/60/EC), adopted in 2007, requires European Union member states to produce preliminary flood risk assessments, flood hazard maps, and flood risk management plans for areas with significant flood risk. The directive specifies three scenarios for the flood hazard maps: low probability (extreme events, including the PMF), medium probability (the 100-year flood, or the "high probability" scenario in the directive's terminology), and high probability (frequent events, typically the 10-year or 30-year flood). The maps must show the flood extent, the water depth, and (where appropriate) the flow velocity. The flood risk maps must show the number of inhabitants affected, the type of economic activity affected, and the installations that might cause accidental pollution. The member states must produce the maps by 2013 (for the flood hazard maps) and review them every 6 years. The maps are made publicly available through the member states' national flood portals and through the European Flood Awareness System. The directive is the

exemplar of the transboundary flood risk mapping: the maps are produced by each member state but follow a common specification, enabling the cross-border coordination of flood risk management.

The flood-prone archipelagic state model is exemplified by Indonesia's Badan Informasi Geospasial (BIG) and the Philippines' National Mapping and Resource Information Authority (NAMRIA). Indonesia, with its 17,000 islands and its location in the Pacific Ring of Fire, faces compound flood hazards: riverine flooding from the monsoon rainfall, coastal flooding from the storm surges, and tsunami-related flooding from the seismic activity. BIG, established in 2011 as the national geospatial information agency reporting directly to the President, produces the national flood hazard maps and the real-time flood monitoring through the InAFLOW (Indonesia Automatic Flood Early Warning System) programme, which integrates the satellite imagery (Sentinel-1 SAR for the flood extent mapping, acquired through the Copernicus programme's free and open data policy), the rain gauge network (operated by the Meteorological, Climatological, and Geophysical Agency, BMKG), and the hydrological models (the HEC-HMS and HEC-RAS models, calibrated against the historical flood events). The InAFLOW programme provides the flood early warning to the National Disaster Management Authority (BNPB) and to the provincial and district emergency management agencies, and the early warning's lead time of 6 to 24 hours allows the evacuation of the populations at risk before the flood arrives. The Philippines, with its 20-25 tropical cyclones per year, faces the most extreme flood hazard of any country in the world. NAMRIA produces the national flood hazard maps and supports the National Disaster Risk Reduction and Management Council's flood response through the satellite-based flood monitoring. The model demonstrates the geospatial intelligence function in jurisdictions with the most demanding flood hazards, where the compound nature of the hazard (riverine + coastal + storm surge) requires the integration of multiple modeling approaches and the continuous monitoring of multiple sensor networks.

These three programmes confirm the domain's doctrine. Flood hazard mapping is the foundation: the NFIP's FIRMs, the Directive's flood hazard maps, and BIG's national flood hazard maps all begin with the modeling of the flood hazard. Exposure assessment is the analytical core: the Directive's flood risk maps, the NFIP's community-level exposure assessments, and BIG's population exposure assessments all overlay the hazard maps with the population, the property, and the infrastructure. Event response is the operational value: the NFIP's flood response support, the Directive's flood risk management plans, and BIG's real-time flood monitoring all provide the intelligence that the emergency management authorities use for the operational response. The three programmes also confirm the tiered classification regime: public flood hazard maps for the public awareness, internal exposure reports for the critical infrastructure protection, and security-confidential scenarios for the dam and levee failure.

19.9 Worked Example: Producing a 100-Year Flood Hazard Map for a Coastal City

The office is tasked with producing a 100-year flood hazard map for Portsmith, a coastal city of 250,000 inhabitants located at the mouth of the Riverbourne. The city faces both riverine flooding (from the Riverbourne) and coastal flooding (from the storm surge), and the 100-year flood hazard map must integrate both hazards.

The hazard specialist begins with the riverine flood hazard. The hydrological model (HEC-HMS) is configured with the 100-year design rainfall (the rainfall depth for the 24-hour, 100-year return period, derived from the national rainfall frequency analysis: 185 millimetres in 24 hours for the Riverbourne catchment). The SCS curve number method is used for the rainfall-runoff transformation, with the curve numbers derived from the land cover and the soil maps (the catchment's weighted curve number is 78, reflecting the mix of urban and agricultural land). The HEC-HMS run produces a peak flow of 1,250 cubic metres per second at the catchment outlet, which is the upstream boundary condition for the hydraulic model.

The coastal flood hazard is derived from the storm surge analysis. The specialist retrieves the 100-year storm surge level from the national storm surge atlas: 2.8 metres above mean sea level for Portsmith. The storm surge level is the downstream boundary condition for the hydraulic model.

The hydraulic model (HEC-RAS 2D) is configured with the LiDAR DEM (2-metre resolution, 0.12-metre vertical RMSE), the standard Manning's n values (0.035 for the main channel, 0.060 for the urban, 0.080 for the forested floodplain, 0.030 for the agricultural floodplain), the upstream boundary condition (the 100-year hydrograph), and the downstream boundary condition (the 100-year storm surge level). The model is run for a 72-hour simulation period (to capture the coincidence of the riverine flood peak and the storm surge peak, which is the worst-case scenario for the coastal city). The model produces the inundation extent and the depth at 2-metre grid resolution.

The specialist validates the model against the observed data. The model's predicted peak stage at the tidal gauge in Portsmith harbour is 3.65 metres (the observed 100-year peak stage from the tidal gauge's 40-year record is 3.52 metres, a difference of 0.13 metres, within the 0.30-metre tolerance). The model's predicted inundation extent is compared with the observed extent from the 2019 flood (which was approximately a 50-year event): the model's extent for the 50-year scenario (run separately) covers 92 per cent of the observed 2019 extent (above the 80 per cent tolerance).

The specialist produces the exposure report. The 100-year inundation extent is overlaid with the cadastre: 1,847 parcels are within the flood hazard area (623 residential, 89 commercial, 12 industrial, 1,123 agricultural). The inundation extent is overlaid with the population data: approximately 18,500 people live within the flood hazard area (7.4 per cent of the city's population). The inundation extent is overlaid with the infrastructure asset registry: 2 substations, 1 water treatment plant, 1 hospital, and 3 emergency facilities are within the flood hazard area.

The specialist produces the evacuation guidance. The inundation extent is overlaid with the road network: 47 kilometres of roads are within the flood hazard area and are impassable during the 100-year flood. The evacuation routes are identified: 3 primary routes remain passable (their elevation is above the 100-year flood level), and 2 shelter locations (a school and a community centre, both outside the flood hazard area, with a combined capacity of 3,000 people) are designated.

The specialist produces the products. The flood hazard map (Public classification) shows the 100-year flood extent and depth, with the blue depth gradient and the parcel boundaries overlaid. The exposure report (Internal classification, because of the critical infrastructure exposure) identifies the 1,847 affected parcels, the 18,500 affected people, and the 2 substations, 1 water treatment plant, 1 hospital, and 3 emergency facilities. The evacuation guidance (Public classification) identifies the 47 kilometres of impassable roads, the 3 primary evacuation routes, and the 2 shelter locations.

This worked example illustrates the flood risk domain's full method: the hydrological modeling (HEC-HMS with the SCS curve number), the coastal flood hazard integration (the storm surge as the downstream boundary condition), the hydraulic modeling (HEC-RAS 2D with the LiDAR DEM and the Manning's n values), the validation (against the tidal gauge and the 2019 event), the exposure assessment (the cadastral, population, and infrastructure overlays), the evacuation guidance (the road network and the shelter analysis), and the classification regime (Public for the hazard map and the evacuation guidance, Internal for the exposure report). The domain's products provide the emergency management authority, the spatial planning authority, and the public with the intelligence they need to prepare for, respond to, and recover from the 100-year flood. The flood hazard map supports the spatial planning authority's zoning decisions and the public's flood awareness. The exposure report supports the emergency management authority's preparedness planning and the sectoral regulators' resilience assessments. The evacuation guidance supports the emergency management authority's evacuation planning and the public's individual preparedness. Together, the three products constitute the office's flood risk intelligence for Portsmith, and the intelligence is the basis for the government's flood risk management decisions for the city. The flood hazard map is incorporated into the city's zoning ordinance, the exposure report informs the emergency management authority's evacuation plan, and the evacuation guidance is distributed to the public through the city's emergency notification system. The office's flood risk intelligence for Portsmith demonstrates the domain's full capability: the integration of riverine and coastal flood hazards, the use of high-resolution LiDAR terrain data, the validation against observed events, the exposure assessment that connects the hazard to the people and the infrastructure, and the classification regime that balances public awareness with critical infrastructure protection.

Chapter 20 — Food Security Intelligence

20.1 The Domain Defined

Food security intelligence is the production of crop condition assessments, yield forecasts, and procurement advisories supporting a strategic reserve authority, such as a national grain reserve or buffer stock scheme. The domain serves the strategic reserve authority as its principal consumer. Its products — the weekly crop condition bulletin, the seasonal yield forecast, and the procurement advisory — support the reserve authority's release and buy decisions. The domain operates on a weekly cadence with seasonal milestones, under a split classification regime: crop condition summaries are Public, while procurement advisories are Internal.

Food security is a matter of national security. The Food and Agriculture Organization of the United Nations estimates that approximately 733 million people faced hunger in 2023, and that global food price volatility — driven by climate variability, geopolitical conflict, supply chain disruption, and economic shocks — has become a persistent feature of the 21st-century food system. Governments that maintain strategic grain reserves (China, with approximately 600 million tonnes of grain reserves; India, with approximately 60 million tonnes; Japan, with approximately 5 million tonnes of rice reserves) use the reserves to stabilize domestic prices, ensure availability during supply disruptions, and protect the population from food insecurity. The geospatial office's food security intelligence supports the reserve authority's management of these reserves by providing the forward-looking intelligence on crop conditions and yield prospects that the authority needs to time its release and buy decisions.

The domain is bounded by three interfaces. The forward-sensor relationship separates it from a crop-monitoring service (Chapter 25): the office produces procurement advisories that synthesize crop condition with reserve stock levels, not merely crop condition assessments. The intelligence-decision division separates it from the reserve authority's decision function: the office produces the advisory; the authority decides. The intelligence-extension boundary separates it from the agricultural extension service: the office produces intelligence; the extension service provides agronomic advice. The boundaries are maintained because the office's value to the reserve authority depends on the office's analytical independence from the authority's procurement decisions, and the office's value to the agricultural ministry depends on the office's independence from the ministry's policy preferences.

20.2 The Scope

The food security domain covers three activities: crop condition assessment, yield forecasting, and procurement advisory production.

Crop condition assessment evaluates the state of crops from satellite imagery, using vegetation indices computed from the calibrated surface reflectance. The assessment is stratified by parcel, using the cadastre feed to identify individual fields, and by crop type, using the farmer-reported planting data transmitted by the agricultural ministry. The assessment is conducted weekly during the growing season, and it provides the reserve authority with a current picture of how the crops are developing across the jurisdiction's agricultural regions. The assessment's key output is the condition score — a

percentile or a departure-from-median that compares the current vegetation index value with the historical distribution for the same parcel, the same crop, and the same phenological stage.

Yield forecasting estimates the crop yield from the condition assessment and the historical data, using crop models that integrate the condition, the weather, and the soil moisture. The yield forecast is produced at seasonal milestones — planting (the initial forecast, based on the planned area and the historical yield), heading (the updated forecast, based on the condition assessment and the weather), and harvest (the final forecast, based on the near-harvest condition and the weather). The yield forecast is the analytical core of the domain, because the forecast is the basis for the procurement advisory's release or buy signal.

The procurement advisory synthesizes the yield forecast with the reserve stock levels to produce a release or buy signal. The advisory is the domain's forward-sensor product — the product that distinguishes the office from a crop-monitoring service. A crop-monitoring service tells the authority how the crops are doing; the procurement advisory tells the authority what to do about it: release grain from the reserve (if the forecast indicates a surplus) or buy grain for the reserve (if the forecast indicates a deficit). The advisory does not make the decision — the authority decides — but the advisory provides the analytical basis for the decision.

The scope does not include agronomic advice, which is the agricultural extension service's function. It does not include market analysis (futures prices, trade flows, commodity speculation), which is the market authority's function. It does not include reserve management (the physical storage, the rotation, the quality control), which is the reserve authority's function. An office that expands its scope to include agronomic advice, market analysis, or reserve management exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

20.3 The Method

The method of the food security domain combines crop modeling with supply-demand analysis. The crop modeling produces the crop condition assessments and the yield forecasts; the supply-demand analysis produces the procurement advisory.

The crop modeling begins with the vegetation indices. The Normalized Difference Vegetation Index (NDVI), proposed by Rouse et al. in 1973, is computed from the calibrated surface reflectance in the red and near-infrared bands: NDVI = (NIR – Red) / (NIR + Red). The NDVI exploits the fact that healthy vegetation absorbs red light for photosynthesis and reflects near-infrared light due to the internal structure of the leaf. The values range from –1 (water) to +1 (dense vegetation), with bare soil at approximately 0.1 to 0.2, sparse vegetation at 0.2 to 0.5, and dense vegetation at 0.6 to 0.8. The NDVI is computed for each parcel from the Sentinel-2 Level-2A surface reflectance, using the red band (Band 4, 665 nanometres, 10-metre resolution) and the near-infrared band (Band 8, 842 nanometres, 10-metre resolution).

The NDVI has limitations. The index saturates at high biomass (above approximately LAI 3.0), meaning that the NDVI cannot distinguish between a healthy crop with LAI 3.0 and a healthy crop with LAI 6.0 — both produce NDVI values near 0.8. The Enhanced Vegetation Index (EVI), proposed by Huete et al. in 2002, addresses this limitation by incorporating the blue band to correct for atmospheric scattering and the canopy background: EVI = G × (NIR – Red) / (NIR + C1 × Red –

C2 × Blue + L), where G = 2.5, C1 = 6, C2 = 7.5, L = 1. The EVI does not saturate as quickly as the NDVI and is more sensitive to high-biomass vegetation. The Normalized Difference Water Index (NDWI), proposed by Gao in 1996, is computed from the near-infrared and the shortwave-infrared bands: NDWI = (NIR – SWIR) / (NIR + SWIR). The NDWI is sensitive to the vegetation's water content and is used to assess the crop's water stress, which is a leading indicator of yield reduction.

The condition score is computed by comparing the current vegetation index value with the historical distribution for the same parcel, the same crop, and the same phenological stage. The phenological stage is determined from the crop calendar (which specifies the expected planting and harvest dates for each crop and each region) and from the accumulated growing degree days (GDD). The GDD is computed as: GDD = max(0, (T_max + T_min) / 2 – T_base), where T_max and T_min are the daily maximum and minimum temperatures, and T_base is the crop's base temperature (the temperature below which the crop does not grow: 0°C for wheat, 10°C for maize, 10°C for rice). The GDD is accumulated from the planting date, and the accumulated GDD determines the crop's phenological stage: for wheat, approximately 500 GDD for tillering, 850 GDD for stem elongation, 1,200 GDD for heading, and 1,700 GDD for maturity.

The yield forecasting combines the condition score with the weather data, the soil moisture data, and the historical yield data. Two approaches are used. The statistical approach uses a regression model that relates the vegetation index at a specific phenological stage to the final yield. The regression is calibrated using historical data (typically 10 to 20 years of yield and vegetation index data for the same parcel or the same district), and the regression coefficients are crop-specific and region-specific. The regression model is simple and robust, but it assumes that the relationship between the vegetation index and the yield is stable over time, which may not hold if the crop variety, the management practice, or the climate changes. The mechanistic approach uses a crop growth model that simulates the crop's physiological processes: photosynthesis, respiration, transpiration, biomass partitioning, and yield formation. The World Food Studies (WOFOST) model, developed by Wageningen University and used by the European Commission's Mars bulletin, and the Decision Support System for Agrotechnology Transfer (DSSAT), developed by the University of Georgia and used by many national agricultural research services, are the principal mechanistic models. The mechanistic model is more complex and requires more data (daily weather, soil properties, crop management) but is more transferable and can simulate the effects of management changes and climate changes on the yield.

The supply-demand analysis synthesizes the yield forecast with the reserve stock levels and the demand assumptions to produce the release or buy signal. The production estimate is computed as: Production = Yield × Planted Area, where the Yield is from the yield forecast and the Planted Area is from the cadastre or from the agricultural ministry's planted area survey. The consumption estimate is computed as: Consumption = Population × Per-Capita Consumption Rate, where the Per-Capita Consumption Rate is derived from the historical consumption data, adjusted for trends (population growth, dietary changes, income growth). The supply-demand balance is: Balance = Production + Reserve Stocks – Consumption. If the Balance is positive (a surplus), the advisory recommends a release from the reserve. If the Balance is negative (a deficit), the advisory recommends a buy for the reserve. If the Balance is near zero, the advisory recommends a hold (no release and no buy).

The method's validation is conducted against the harvest outcomes. After the harvest, the agricultural ministry's statistical service publishes the reported yield (the actual yield, as reported by the farmers and verified by the statistical service's field surveys). The validation compares the forecast yield with the reported yield: if the forecast is within 10 per cent of the reported yield at the district level, the forecast is accurate. The validation results are recorded in the provenance chain, and the regression coefficients or the crop model parameters are recalibrated based on the validation to improve the next season's forecast.

20.4 The Products

The food security domain produces three products, each serving a different phase of the reserve management cycle.

The weekly crop condition bulletin is produced every week during the growing season and delivered by close of business each Wednesday. The bulletin comprises three sections. The crop condition assessment section presents the condition map (the green-to-red condition gradient, with the parcel boundaries overlaid), the condition statistics (the percentage of the planted area in each condition class: above-normal, normal, below-normal, poor, very poor), and the condition trend (the change from the previous week's assessment). The yield forecast update section presents the updated yield forecast (by district and by region, with the forecast's confidence interval), the forecast's assumptions (the weather scenario, the soil moisture, the management conditions), and the forecast's change from the previous update. The seasonal progress summary section presents the phenological stage (the accumulated GDD, the expected stage transition dates), the weather summary (the seasonal rainfall, the temperature, the comparison with the normal), and the seasonal outlook (the remaining growing period, the expected weather, the risk factors). The bulletin's timeliness and accuracy are the conditions of the reserve authority's trust.

The seasonal yield forecast is produced at three seasonal milestones. The planting forecast (produced at the beginning of the growing season) is based on the planned area, the historical yield, and the seasonal weather outlook. The planting forecast is the initial forecast, and its accuracy is limited by the uncertainty in the weather outlook. The heading forecast (produced at the heading stage, approximately mid-season) is based on the condition assessment, the weather to date, and the weather outlook for the remaining growing period. The heading forecast is more accurate than the planting forecast, because the condition assessment reflects the actual crop conditions rather than the assumptions. The harvest forecast (produced approximately two weeks before the expected harvest) is based on the near-harvest condition, the weather to date, and the expected harvest conditions. The harvest forecast is the most accurate forecast, because the crop is near maturity and the weather's remaining influence is limited. The seasonal yield forecast provides the authoritative yield estimate for the reserve authority's seasonal planning.

The procurement advisory synthesizes the yield forecast with the reserve stock levels to produce the release or buy signal. The advisory comprises four sections. The yield forecast section presents the current yield forecast and its change from the previous forecast. The stock level section presents the current reserve stock levels (obtained from the reserve authority through the data-sharing arrangement) and the stock-to-consumption ratio (the number of months of consumption that the stocks represent, which is the standard indicator of the reserve's adequacy). The supply-demand

analysis section presents the production estimate, the consumption estimate, and the supply-demand balance. The signal section presents the advisory's conclusion: release (with the recommended release quantity), buy (with the recommended buy quantity), or hold (with the rationale for the hold). The advisory is the domain's forward-sensor product, produced as the yield forecast and the stock levels warrant, and delivered to the reserve authority through the internal distribution channel.

20.5 The Consumers

The food security domain serves three consumers, each with different needs and different decision contexts.

The strategic reserve authority is the principal consumer. The authority uses the procurement advisory for the release and buy decisions (the authority's operational decisions), the yield forecast for the seasonal planning (the authority's medium-term planning), and the crop condition bulletin for the ongoing awareness (the authority's situational awareness). The authority's release decision is: should the authority release grain from the reserve to the market, stabilizing the price and increasing the availability? The authority's buy decision is: should the authority buy grain from the market for the reserve, building the stock for future needs? The authority makes the decisions, not the office, but the authority's decisions are informed by the office's advisory. The authority requires the advisory in a format that supports the decision (a one-page summary with the signal, the rationale, and the confidence level) and on a timeline that allows the decision to be made before the market conditions change.

The agricultural ministry is a secondary consumer. The ministry uses the crop condition bulletin for the agricultural policy (the ministry's assessment of the sector's performance, the identification of regions requiring support, the planning of the extension services) and the yield forecast for the production estimates (the ministry's contribution to the national economic statistics, the food balance sheets, and the trade policy). The ministry requires the bulletin in a format that supports the policy analysis (the condition statistics by district, the trend analysis, the comparison with the historical conditions) and on the weekly cadence that matches the ministry's reporting cycle.

The public is a tertiary consumer. The public uses the public crop condition summary for the food security awareness (the public's understanding of the current food supply situation, which affects the public's confidence in the food supply and the public's purchasing behaviour). The public requires the summary in an accessible format (a web map, a press release, or a social media post) and in a language that the public understands. The public summary does not include the procurement advisory or the detailed yield forecast, because the premature disclosure of the advisory or the forecast could move market prices.

20.6 The Cadence

The food security domain operates on a weekly cadence with seasonal milestones. The weekly cadence is the domain's operational rhythm, and it is maintained throughout the growing season (typically 4 to 6 months, depending on the crop and the climate). The weekly cadence comprises: the acquisition of the latest satellite imagery (Sentinel-2, with a 5-day revisit cycle, meaning that a new image is available for each parcel approximately every 5 days, subject to cloud cover), the processing of the imagery (the surface reflectance calibration, the cloud masking, the NDVI/EVI/NDWI

computation), the condition assessment (the comparison with the historical distribution, the condition score computation, the parcel-level stratification), the yield forecast update (the integration of the condition assessment with the weather data and the historical yield, the regression or the mechanistic model run), and the production and delivery of the crop condition bulletin by close of business each Wednesday.

The seasonal milestones are the domain's seasonal rhythm, and they are maintained at three points in the growing season. The planting milestone (the beginning of the growing season) comprises the initial yield forecast based on the planned area and the historical yield, and the planting confirmation (the verification that the planned area has been planted, using the satellite imagery to detect the crop emergence). The heading milestone (the mid-season) comprises the updated yield forecast based on the condition assessment and the weather to date, and the condition assessment's seasonal review (the comparison of the current season's condition with the historical conditions at the same stage). The harvest milestone (the end of the growing season) comprises the final yield forecast based on the near-harvest condition, the harvest monitoring (the satellite-based detection of the harvest activity, using the change in the vegetation index from green to bare soil), and the post-harvest validation (the comparison of the forecast yield with the reported harvest yield).

The weekly cadence is maintained by the resource economics specialist, who produces the bulletin, and the production operator, who produces and delivers it. The seasonal milestones are maintained by the resource economics specialist, who produces the yield forecasts, and the technical lead, who validates them. During the growing season, the weekly cadence intensifies as the harvest approaches: the bulletin's content expands to cover the changing crop conditions, the specialist's modeling runs become more frequent (from weekly to daily in the final two weeks before the harvest), and the procurement advisory's production shifts from monthly to weekly as the yield forecast's accuracy improves.

Outside the growing season (the off-season), the weekly cadence is suspended, and the domain operates on a monthly cadence: the specialist monitors the off-season conditions (the soil moisture, the fallow land management, the off-season rainfall) and produces a monthly off-season summary. The off-season monitoring is less intensive than the growing-season monitoring, but it provides the reserve authority with the information needed to assess the conditions for the next growing season.

20.7 The Classification Regime

The food security domain operates under a split classification regime. The split regime reflects the dual character of the domain's intelligence: the crop condition information serves the public awareness, and the procurement advisory serves the reserve authority's procurement position.

Crop condition summaries are Public. Public awareness of the crop condition supports public confidence in the food supply, and public confidence is the condition of the social stability that the reserve authority's price stabilization seeks to maintain. If the public knows that the crop condition is good, the public's confidence in the food supply is high, and the market prices remain stable. If the public does not know the crop condition, the public's uncertainty may lead to panic buying and price speculation, undermining the reserve authority's price stabilization efforts. The United States Department of Agriculture's World Agricultural Production circular, which publishes the crop condition assessments and the yield forecasts for the major producing countries, is publicly available,

and the European Commission's Mars bulletin is also publicly available. The Public classification of the crop condition summary is consistent with the comparative practice.

Procurement advisories are Internal. The premature disclosure of the advisory could move market prices, because the advisory's release or buy signal indicates the reserve authority's intended market action, and the market participants who obtain the signal before the authority acts can front-run the authority's action (buying grain before the authority's buy, driving up the price; or selling grain before the authority's release, driving down the price). The Internal marking restricts the advisory to the reserve authority and the office, and the internal distribution channel (the secure data interface between the office and the authority, with access logging and access review) enforces the restriction. The Internal classification of the procurement advisory is consistent with the practice of national reserve authorities, which treat their procurement intentions as confidential market-sensitive information.

The split regime is maintained by the production operator, who applies the markings, and the office manager, who reviews the markings before release. A product that is misclassified has compromised either public awareness (if the crop condition summary is over-classified as Internal, restricting it from the public) or the reserve authority's procurement position (if the procurement advisory is under-classified as Public, disclosing it to the market). The regime allows the domain to serve public awareness and the reserve authority's procurement simultaneously, and the simultaneous service is the condition of the domain's contribution to the jurisdiction's food security.

20.8 The Comparative Practice

The GEOGLAM initiative — the Group on Earth Observations Global Agricultural Monitoring — was launched in 2011 by the G20 Agricultural Ministers in response to the 2007-2008 global food price crisis, which demonstrated that the lack of timely, reliable, and transparent crop condition information contributed to the panic-driven price spikes that caused food insecurity in import-dependent countries. GEOGLAM's Crop Monitor provides a monthly global crop condition assessment that synthesizes the assessments of 35 participating organizations (including the FAO, the USDA, the European Commission's JRC, the Chinese Academy of Sciences, the Indian Space Research Organisation, and the national agricultural monitoring services of the major producing countries). The Crop Monitor uses a common condition classification (favorable, watch, poor, failure) and a common reporting format, enabling the cross-comparison of the conditions across countries and regions. GEOGLAM is the exemplar of the international coordination of crop monitoring, and its monthly assessment is the global benchmark for the crop condition information.

The European Commission's Mars (Monitoring Agricultural Resources) bulletin, produced by the Mars Unit of the Joint Research Centre since 1992, is the longest-running operational crop monitoring service. The Mars bulletin provides crop condition assessments and yield forecasts for the European Union's 27 member states and for approximately 30 other countries, using the WOFOST crop growth model (calibrated for each crop and each region), the MARS (Monitoring Agriculture with Remote Sensing) database of weather stations (approximately 5,000 stations across Europe and the Mediterranean), and the satellite imagery from Sentinel-2 and Proba-V. The Mars bulletin is published monthly during the growing season and is publicly available through the JRC's Mars bulletin website. The bulletin's yield forecasts are used by the European Commission's Directorate-

General for Agriculture for the Common Agricultural Policy's market management and by the member states' agricultural ministries for their national food security planning.

The United States Department of Agriculture's Foreign Agricultural Service (FAS) produces the World Agricultural Production (WAP) circular, a monthly publication that provides crop condition assessments and yield forecasts for the major producing countries. The FAS uses a combination of satellite imagery analysis (from the Global Agricultural Monitoring (GLAM) system, which provides near-real-time NDVI data from MODIS and VIIRS), crop models (the Crop Assessment Tool, which integrates the NDVI, the weather, and the historical yield), and agricultural attaché reports (the FAS's overseas officers, stationed in approximately 100 countries, who report on the local crop conditions and the agricultural policy). The WAP circular is publicly available through the FAS website, and its yield forecasts are the global benchmark for the international grain trade.

The national crop monitoring services — China's Institute of Agricultural Resources and Regional Planning (which uses the FY-3 satellite's MERSI sensor and the HJ-1 satellite's CCD sensor for the crop condition monitoring), India's Mahalanobis National Crop Forecast Centre (which uses the Resourcesat-2 satellite's AWiFS sensor for the crop area estimation and the RISAT-1 satellite's SAR for the crop condition monitoring during the monsoon cloud cover), and Brazil's National Supply Company (which uses the CBERS-4 satellite's MUX sensor and the Landsat imagery for the crop area estimation and the crop condition monitoring) — are the exemplars of the national crop monitoring, and their products are the products that the geospatial office's products complement.

These programmes confirm the domain's doctrine. Crop condition assessment is the foundation: GEOGLAM's Crop Monitor, the Mars bulletin, and the FAS WAP all begin with the satellite-based vegetation index analysis. Yield forecasting is the analytical core: the Mars bulletin's WOFOST model, the FAS's Crop Assessment Tool, and the national services' regression and mechanistic models all produce the yield forecast that the reserve authorities and the markets use. The procurement advisory is the forward-sensor value: the advisory synthesizes the yield forecast with the reserve stock levels, and the synthesis is the product that distinguishes the geospatial intelligence office from a crop-monitoring service. The programmes also confirm the split classification regime: the crop condition assessments are public (the Mars bulletin, the WAP circular, and the GEOGLAM Crop Monitor are all publicly available), while the procurement intentions of the reserve authorities are confidential.

20.9 Worked Example: Producing a Procurement Advisory During a Drought Year

The resource economics specialist is producing the procurement advisory for the week of 15 July 2024. The growing season has been marked by a developing drought: the seasonal rainfall (March to July) is 35 per cent below normal for the jurisdiction's principal agricultural region, the Northfields, which produces approximately 60 per cent of the jurisdiction's wheat. The specialist has been tracking the drought's impact on the crop condition through the weekly bulletins, and the condition has deteriorated from "above-normal" in March (when the planting was completed under favourable conditions) to "below-normal" in May (when the rainfall deficit began to affect the crop) to "poor" in July (when the drought intensified and the crop's yield potential was significantly reduced).

The specialist retrieves the latest Sentinel-2 image (acquired 13 July 2024, 8 per cent cloud cover) and processes it through the crop model plug-in. The NDVI for the Northfields' wheat parcels is at the 22nd percentile of the historical distribution (compared with the 62nd percentile at the same stage in 2023, a normal year). The EVI is at the 25th percentile, and the NDWI is at the 18th percentile, confirming the crop's water stress. The condition score for the Northfields' wheat is "poor," with 45 per cent of the planted area in the poor class, 30 per cent in the below-normal class, 20 per cent in the normal class, and 5 per cent in the above-normal class.

The specialist updates the yield forecast. The statistical regression model (calibrated with 15 years of yield and NDVI data for the Northfields' wheat) predicts a yield of 2.7 tonnes per hectare, compared with the 10-year average of 4.2 tonnes per hectare (a 36 per cent reduction) and the 2023 yield of 4.5 tonnes per hectare (a 40 per cent reduction). The mechanistic WOFOST model (configured with the current weather data and the soil moisture data) predicts a yield of 2.9 tonnes per hectare. The specialist uses the average of the two models (2.8 tonnes per hectare) as the forecast yield, with a confidence interval of ±0.4 tonnes per hectare.

The specialist computes the production estimate. The planted area for the Northfields' wheat is 850,000 hectares (from the cadastre feed, verified against the agricultural ministry's planted area survey). The production estimate is: 2.8 tonnes/hectare × 850,000 hectares = 2,380,000 tonnes. The specialist also computes the production for the jurisdiction's other agricultural regions (which are less affected by the drought): the Southfields' wheat (400,000 hectares, yield 3.8 tonnes/hectare, production 1,520,000 tonnes) and the Westfields' wheat (250,000 hectares, yield 4.0 tonnes/hectare, production 1,000,000 tonnes). The total wheat production estimate is: 2,380,000 + 1,520,000 + 1,000,000 = 4,900,000 tonnes.

The specialist retrieves the reserve stock levels from the strategic reserve authority. The current wheat stocks are 1,800,000 tonnes (obtained through the data-sharing arrangement with the authority). The stock-to-consumption ratio is: 1,800,000 / (150,000 tonnes/month) = 12 months of consumption.

The specialist computes the supply-demand balance. The annual wheat consumption is: 150,000 tonnes/month × 12 months = 1,800,000 tonnes. The balance is: 4,900,000 (production) + 1,800,000 (reserve stocks) – 1,800,000 (consumption) = 4,900,000 tonnes. However, this balance does not account for the carry-over (the amount of wheat that must be retained in the reserve for the next year's food security). The minimum carry-over is 6 months of consumption (900,000 tonnes, the reserve authority's minimum stock threshold). The adjusted balance is: 4,900,000 + 1,800,000 – 1,800,000 – 900,000 = 4,000,000 tonnes. The positive balance indicates a surplus, but the surplus is driven by the reserve stocks (1,800,000 tonnes), not by the production (4,900,000 tonnes, which is 22 per cent below the 5-year average of 6,300,000 tonnes).

The specialist produces the procurement advisory. The advisory's signal is "hold" (not release and not buy): the current reserve stocks (1,800,000 tonnes, 12 months of consumption) are adequate, but the production (4,900,000 tonnes) is significantly below the 5-year average, and a release would deplete the reserve to a level that may not be adequate for the next year if the drought persists. The advisory recommends that the authority hold the current stocks, monitor the crop condition in the remaining growing period (approximately 6 weeks until harvest), and reassess the procurement decision at the harvest milestone when the final yield forecast is available.

The specialist applies the Internal classification marking to the advisory (because the premature disclosure of the hold signal and the production estimate could move market prices) and submits the advisory to the production operator for production. The production operator produces the advisory, the office manager reviews it (verifying the classification marking, the analytical content, and the signal), and the advisory is delivered to the strategic reserve authority through the internal distribution channel.

This worked example illustrates the food security domain's full method: the vegetation index analysis (NDVI at the 22nd percentile, EVI at the 25th percentile, NDWI at the 18th percentile), the yield forecasting (the statistical regression at 2.7 t/ha and the WOFOST at 2.9 t/ha, averaged to 2.8 t/ha with a confidence interval), the production estimation (2.8 t/ha × 850,000 ha = 2,380,000 tonnes for the Northfields, 4,900,000 tonnes total), the supply-demand analysis (production + stocks – consumption – carry-over = 4,000,000 tonnes surplus, driven by stocks not production), and the procurement advisory's signal (hold, with the rationale that a release would deplete the reserve below the adequate level for the next year if the drought persists). The domain's products provide the reserve authority with the intelligence it needs to make the procurement decision that maintains the jurisdiction's food security during a drought year, and the intelligence is the forward-sensor value that distinguishes the geospatial office from a crop-monitoring service.

Chapter 21 — Energy Infrastructure Intelligence

21.1 The Domain Defined

Energy infrastructure intelligence is the monitoring of a jurisdiction's energy infrastructure — generation facilities, transmission corridors, substations, and pipelines — for encroachment, degradation, and vulnerability. Energy infrastructure is critical infrastructure: its disruption affects the economy, the public safety, and the national security. The 2003 Northeast Blackout (United States and Canada, August 2003) affected 55 million people and caused approximately $6 billion in economic damage, triggered by a cascading failure that began with a transmission line sagging into overgrown vegetation in Ohio. The 2017 Hurricane Harvey (United States, August 2017) disrupted approximately 25 per cent of the United States' refining capacity and caused fuel price spikes across the country. The 2021 Colonial Pipeline cyberattack (United States, May 2021) caused fuel shortages across the eastern United States for six days. These events demonstrate that the energy infrastructure's vulnerability to physical encroachment, environmental degradation, and malicious attack is a matter of national security, and the geospatial office's intelligence supports the infrastructure's protection.

The domain serves the energy sector entity and the sectoral regulator as its principal consumers. Its products — the infrastructure risk map, the encroachment alerts, and the maintenance advisories — support the entity's asset management and the regulator's compliance oversight. The domain operates on a monthly cadence with event-driven response, under a Security Confidential classification by default: the geolocation of critical energy infrastructure is security-sensitive, and the products that identify the infrastructure and its vulnerabilities are restricted to the security-cleared personnel.

The domain is bounded by three interfaces. The intelligence-regulation boundary separates it from the sectoral regulator's regulatory function (Chapter 26): the geospatial office produces the intelligence, and the regulator regulates the sector. The intelligence-operations boundary separates it from the energy sector entity's operational function: the geospatial office monitors the infrastructure, and the entity operates it. The intelligence-response division separates it from the emergency management authority's response function (Chapter 24): the geospatial office produces the intelligence, and the authority responds to the incidents.

21.2 The Scope

The energy infrastructure domain covers three activities: infrastructure monitoring, encroachment detection, and vulnerability assessment.

Infrastructure monitoring is the regular monitoring of the energy infrastructure through high-resolution satellite imagery. The monitoring covers four infrastructure types. Generation facilities (power plants, refineries, LNG terminals) are monitored for physical changes: new construction within the facility's perimeter, changes to the storage areas, and changes to the access roads. Transmission corridors (high-voltage transmission lines, pipeline rights-of-way) are monitored for encroachment: vegetation growth into the clearance zone, third-party construction within the right-of-way, and ground subsidence along the corridor. Substations (the switching and transformation facilities that connect the transmission and distribution networks) are monitored for physical

changes: new construction within the substation's perimeter, changes to the perimeter security, and changes to the access roads. Pipelines (the oil and gas pipelines that transport the fuel) are monitored for ground disturbance: excavation within the right-of-way (which may indicate third-party construction or, in some jurisdictions, illegal tapping), subsidence along the pipeline route (which may stress the pipe), and vegetation changes along the route (which may indicate leaks).

Encroachment detection is the detection of the third-party construction, the vegetation growth, and the other changes that may threaten the infrastructure. The detection is conducted through the change detection algorithms of the change detection plug-in (Chapter 10, section 10.4), which compare the current month's imagery with the previous month's imagery and identify the pixels or objects that have changed. The detected changes are classified by type: vegetation encroachment (the vegetation has grown into the clearance zone), third-party construction (a new structure has been built within the right-of-way), ground subsidence (the ground elevation has decreased, indicating potential slope instability or underground voids), and excavation (the ground surface has been disturbed, indicating potential damage to buried infrastructure).

Vulnerability assessment is the assessment of the infrastructure's vulnerability to the detected changes, conducted through the vulnerability scoring algorithms that prioritize the changes for the consumer. The vulnerability score combines the asset's criticality (a transmission line serving 50,000 customers is more critical than one serving 5,000), the change's proximity (a change within 5 metres of a transmission tower is more urgent than a change at 25 metres), and the change's type (ground subsidence near a pipeline is more urgent than vegetation growth near a fence line).

The scope does not include the infrastructure operation (the generation, the transmission, the distribution of energy), which is the energy sector entity's function. It does not include the infrastructure regulation (the safety standards, the reliability standards, the tariff regulation), which is the sectoral regulator's function. It does not include the incident response (the emergency repair, the power restoration, the evacuation), which is the emergency management authority's function. An office that expands its scope to include the operation, the regulation, or the response exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

21.3 The Method

The method of the energy infrastructure domain is satellite-based change detection, supplemented by vulnerability scoring and validated against ground truth inspections.

The change detection compares two acquisitions of the same area and identifies the pixels or objects that have changed. The office uses high-resolution commercial satellite imagery for the change detection, because the moderate-resolution free and open imagery (Sentinel-2 at 10 metres, Landsat at 30 metres) is insufficient to detect the encroachment on the energy infrastructure, which often involves changes at the sub-metre scale: a new building within a substation's perimeter, a vegetation incursion into a 5-metre clearance zone, or a ground subsidence of 0.5 metres along a pipeline route. The office tasks Maxar WorldView-3 (0.31-metre panchromatic, 1.24-metre multispectral) or Airbus Pleiades Neo (0.3-metre panchromatic, 1.2-metre multispectral) for the monthly monitoring, with the tasking parameters specifying the infrastructure corridors as the areas of interest, the 10 per cent maximum cloud cover, and the monthly acquisition window.

Two change detection approaches are used. The image-to-image approach compares the radiometric values of the two acquisitions: the image differencing subtracts the earlier image's pixel values from the later image's pixel values and identifies the pixels where the difference exceeds a threshold. The approach is simple and computationally efficient, but it is sensitive to radiometric differences between the acquisitions (different sun angles, different atmospheric conditions, different sensor calibrations) and to geometric misregistration (if the two images are not precisely co-registered, every pixel appears to have changed). The object-to-object approach first segments the imagery into objects (contiguous regions of similar spectral characteristics, produced by the image segmentation algorithm) and then compares the objects between acquisitions: an object that has changed in area, shape, or spectral characteristics is identified as a change. The object-based approach is more robust to radiometric differences and geometric misregistration, because the comparison is at the object level rather than the pixel level, but it requires the segmentation step, which adds complexity and processing time.

The encroachment detection applies the change detection to a buffer around each asset in the infrastructure asset registry. The buffer width is determined by the asset's type and the threat's characteristics: a transmission line may have a buffer of 30 metres (the vegetation clearance zone specified in the regulation), a substation may have a buffer of 100 metres (the security perimeter), a pipeline may have a buffer of 15 metres (the right-of-way). The detected changes within the buffer are classified by type: vegetation encroachment (the vegetation has grown into the buffer, identified by the increase in the NDVI within the buffer), third-party construction (a new structure has appeared within the buffer, identified by the change from a non-urban to an urban spectral signature), ground subsidence (the ground elevation has decreased within the buffer, identified by the comparison of the DEMs from two acquisitions or by the InSAR analysis of the Sentinel-1 radar data — Interferometric Synthetic Aperture Radar, which measures the ground deformation by comparing the phase difference between two radar acquisitions of the same area, and can detect subsidence at the millimetre scale over periods of weeks to months, making it the preferred method for monitoring the ground stability along the pipeline routes and near the transmission towers), and excavation (the ground surface has been disturbed within the buffer, identified by the change from a vegetated to a bare-soil spectral signature).

The vulnerability scoring combines the detected changes with the asset's criticality and the change's proximity to produce a score that prioritizes the changes for the consumer. The asset's criticality is determined by the asset registry: the registry assigns each asset a criticality rating (critical, high, moderate, low) based on the asset's function (a transmission interconnection is critical; a distribution feeder is moderate), the asset's capacity (a 500 kV transmission line is critical; a 33 kV distribution line is moderate), and the asset's redundancy (a transmission line with no alternative route is critical; a transmission line with a parallel route is high). The change's proximity is the distance from the change to the asset, measured in the GIS. The vulnerability score is computed as: Score = Criticality Weight × Proximity Factor × Change Type Factor, where the Criticality Weight is 4 for critical, 3 for high, 2 for moderate, 1 for low; the Proximity Factor is 1.0 for changes within 5 metres, 0.7 for 5-15 metres, 0.4 for 15-30 metres, 0.1 for beyond 30 metres; and the Change Type Factor is 1.0 for excavation and subsidence, 0.7 for third-party construction, 0.5 for vegetation encroachment. The

score ranges from 0 to 4, and the scores above 2.0 are flagged as high-priority for the consumer's maintenance dispatch.

The method's validation is conducted against the ground truth inspections provided by the energy sector entity's field teams. The field teams inspect the infrastructure corridors on a regular basis (typically quarterly) and report the conditions they observe: the vegetation height, the encroachment structures, the ground conditions, and the infrastructure's physical state. The validation compares the change detection's results with the field teams' reports: the true positives (changes detected by the algorithm and confirmed by the field team), the false positives (changes detected by the algorithm but not confirmed by the field team), and the false negatives (changes detected by the field team but missed by the algorithm). The accuracy metrics (the detection rate and the false alarm rate) are calculated and recorded in the provenance chain, and the algorithm's parameters (the change threshold, the buffer widths, the vulnerability scoring weights) are adjusted to optimize the accuracy.

21.4 The Products

The energy infrastructure domain produces four product types, each serving a different phase of the infrastructure protection cycle.

The infrastructure risk map is the domain's standing product, updated on a monthly cadence. The map shows the monitored infrastructure (the generation facilities, the transmission corridors, the substations, and the pipelines, overlaid on the high-resolution satellite imagery) and the vulnerability scores (colour-coded by priority: red for high-priority, orange for moderate-priority, yellow for low-priority). The map is used by the energy sector entity for the infrastructure management (the entity's maintenance planning, the resource allocation, the risk assessment) and by the sectoral regulator for the regulatory oversight (the regulator's compliance verification, the enforcement actions, the performance assessment). The map carries the Security Confidential classification, because it identifies the critical energy infrastructure's geolocation and its vulnerabilities.

The encroachment alerts are the domain's event-driven product, produced as the encroachments are detected. The alerts identify the specific encroachment (the type, the location, the proximity to the asset, the vulnerability score) and provide the supporting imagery (the before-and-after image pair, annotated to show the encroachment). The alerts are delivered to the energy sector entity's operations centre through the secure data interface, and the entity's maintenance dispatch uses the alerts to prioritize the field response. The alerts carry the Security Confidential classification.

The maintenance advisories are the domain's analytical product, produced for the energy sector entity's maintenance planning. The advisories identify the infrastructure at risk from the ground conditions (the subsidence, the erosion, the slope instability) and recommend the maintenance actions (the vegetation clearing, the ground stabilization, the structural reinforcement). The advisories are produced quarterly, synthesizing the monthly monitoring results and the ground condition analysis into a longer-term planning document. The advisories carry the Security Confidential classification.

Event-specific intelligence is produced when an incident is reported or detected. The intelligence comprises the event-specific infrastructure risk map (showing the incident's location and the affected infrastructure), the event-specific encroachment alerts (identifying any encroachments that may have contributed to the incident), and the event-specific vulnerability assessment (assessing

the incident's impact on the infrastructure's overall vulnerability). The event-specific intelligence is used by the energy sector entity and the emergency management authority for the incident response. The intelligence carries the Security Confidential classification.

21.5 The Consumers

The energy infrastructure domain serves three consumers, each with different needs and different decision contexts.

The energy sector entity is the principal consumer. The entity uses the infrastructure risk map for the infrastructure management (the entity's monthly maintenance planning, the resource allocation, the risk assessment), the encroachment alerts for the encroachment response (the entity's dispatch of the maintenance crews to the encroachment sites), and the maintenance advisories for the maintenance planning (the entity's quarterly maintenance budgeting, the capital investment planning, the asset replacement scheduling). The entity requires the intelligence in operational formats (GeoTIFF for its GIS system, Shapefile for its asset management system) and on the monthly cadence that matches the entity's maintenance cycle.

The sectoral regulator is a secondary consumer. The regulator uses the infrastructure risk map for the regulatory oversight (the regulator's verification that the entity is maintaining the infrastructure in accordance with the safety and reliability standards), the encroachment alerts for the regulatory enforcement (the regulator's verification that the entity is responding to the encroachments within the required timeframe), and the maintenance advisories for the compliance assessment (the regulator's assessment of the entity's maintenance programme's adequacy). The regulator requires the intelligence in regulatory formats (compliance reports, audit summaries) and on the quarterly cadence that matches the regulator's compliance cycle. The regulator's compliance verification checks whether the energy sector entity is maintaining the vegetation clearance zones along the transmission corridors, whether the entity is responding to the encroachment alerts within the required timeframe (typically 30 days for high-priority encroachments), and whether the entity's maintenance programme is adequate to prevent the infrastructure's degradation. The regulator's enforcement actions may include fines, compliance orders, or, in extreme cases, the revocation of the entity's operating licence, and the regulator's decisions are informed by the geospatial office's intelligence. The regulator's relationship with the geospatial office is governed by the intelligence-regulation boundary (Chapter 26): the office produces the intelligence, and the regulator regulates the sector, and the boundary ensures that the office's intelligence is not distorted by the regulator's enforcement priorities.

The emergency management authority is a tertiary consumer. The authority uses the event-specific intelligence for the incident response: when an energy infrastructure incident occurs (a transmission line failure, a pipeline leak, a substation fire), the authority needs the event-specific intelligence to assess the incident's impact on the public safety (the population affected by the power outage, the area affected by the leak, the evacuation zones). The authority requires the intelligence on event-driven timelines (within hours of the incident) and in the operational formats that the authority's emergency operations centre uses.

21.6 The Cadence

The energy infrastructure domain operates on a monthly cadence with event-driven response. The monthly cadence comprises the acquisition of the high-resolution satellite imagery (tasked at the beginning of each month, with the acquisition window of 1-10 days, subject to cloud cover), the processing of the imagery (the georeferencing, the calibration, the mosaicking, and the cloud masking, conducted by the data pipeline operator), the change detection run (conducted by the hazard specialist using the change detection plug-in, typically taking 1-2 days for a full monitoring cycle covering all the infrastructure corridors), the review of the detected changes (the specialist's review of each detected change for plausibility, distinguishing the true changes from the false alarms), the cross-referencing with the infrastructure asset registry (the identification of the affected assets and the computation of the vulnerability scores), and the production and delivery of the infrastructure risk map, the encroachment alerts, and the maintenance advisories by the fifth business day of each month.

The event-driven response is activated when an incident is reported or detected. The incident may be reported by the energy sector entity (a transmission line failure, a pipeline leak), detected by the office's monitoring (a significant ground subsidence along a pipeline route, detected in the monthly change detection run), or reported by a third party (a public report of a vegetation fire near a substation). The event-driven response comprises the acquisition of the event-specific imagery (the office tasks a new satellite acquisition of the incident area, with the shortest possible acquisition window), the processing of the imagery (the expedited processing, with the validation waived only if the office manager explicitly authorizes the waiver), the change detection run (the comparison of the event-specific imagery with the most recent monthly imagery), and the production and delivery of the event-specific intelligence to the energy sector entity and the emergency management authority.

The monthly cadence is maintained by the hazard specialist, who conducts the change detection and the vulnerability assessment, and the production operator, who produces and delivers the monthly products. The event-driven response is maintained by the hazard specialist, who produces the event-specific intelligence, and the office manager, who activates the response and coordinates with the energy sector entity and the emergency management authority.

21.7 The Classification Regime

The energy infrastructure domain operates under a Security Confidential classification by default. The geolocation of the critical energy infrastructure is security-sensitive, because the disclosure of the geolocation could inform attacks on the infrastructure. A terrorist group or a hostile state that obtains the precise geolocation of the critical substations, the transmission interconnections, and the pipeline junctions can plan attacks that maximize the disruption to the energy system, and the attacks' cascading effects on the economy, the public safety, and the national security could be catastrophic. The Security Confidential classification restricts the domain's products to the security-cleared personnel: the energy sector entity's security-cleared staff, the sectoral regulator's security-cleared staff, the emergency management authority's security-cleared staff, and the intelligence agencies.

The Security Confidential classification is consistent with the critical infrastructure information protection frameworks in most jurisdictions. The United States' Critical Infrastructure Information Act of 2002 (Title II of the Homeland Security Act, 6 U.S.C. § 133) establishes the Protected Critical Infrastructure Information (PCII) programme, which protects the critical infrastructure information that is voluntarily submitted to the Department of Homeland Security from public disclosure. The European Union's Directive 2008/114/EC on the identification and designation of European critical infrastructures requires the member states to protect the sensitive information about the critical infrastructures. The United Kingdom's Centre for the Protection of National Infrastructure (CPNI) classifies the critical infrastructure information as "restricted" or "confidential" depending on the sensitivity. The geospatial office's Security Confidential classification is aligned with these frameworks.

The classification regime is maintained by the production operator, who applies the markings, and the office manager, who reviews the markings before release. The regime is non-negotiable: every product carries the Security Confidential marking, and the marking is applied in the standard format (capital letters, top and bottom of every page, in the metadata) and in the standard location. The office manager may declassify a product (reduce the classification from Security Confidential to Internal or Public) only if the product's content does not include the critical infrastructure's geolocation or vulnerabilities, and the declassification is documented in the production log with the rationale. The declassification is rare, because most of the domain's products include the infrastructure's geolocation by their nature (the infrastructure risk map shows the infrastructure's location; the encroachment alerts identify the encroachment's location relative to the infrastructure).

21.8 The Comparative Practice

The United States' critical infrastructure protection programme, established by the Homeland Security Presidential Directive 7 (HSPD-7) in 2003 and updated by Presidential Policy Directive 21 (PPD-21) in 2013, designates the energy sector as one of the 16 critical infrastructure sectors and assigns the Department of Energy as the sector-specific agency. The Department of Energy's Office of Cybersecurity, Energy Security, and Emergency Response (CESER) coordinates the energy sector's physical and cybersecurity, including the satellite-based monitoring of the critical energy infrastructure. The Department of Homeland Security's National Protection and Programs Directorate (NPPD, renamed the Cybersecurity and Infrastructure Security Agency, CISA, in 2018) provides the cross-sector coordination, including the satellite imagery acquisition through the National Geospatial-Intelligence Agency's EnhancedView programme. The programme's products are classified under the PCII framework and are restricted to the authorized federal, state, and private-sector personnel.

The European Union's Copernicus Emergency Management Service (CEMS), operated by the European Commission's Joint Research Centre, provides the satellite-based monitoring of the critical infrastructure during emergencies. The CEMS Rapid Mapping component provides the satellite-based damage assessment within hours of an incident, using the Sentinel imagery and the commercial high-resolution imagery. The CEMS is activated by the authorised users (the national civil protection authorities and the European Commission's Emergency Response Coordination Centre) and provides the products through the Copernicus Emergency Management Service portal. The CEMS's products

are not classified by default (the Sentinel imagery is free and open), but the products that include the critical infrastructure's geolocation are restricted under the Directive 2008/114/EC.

The European Maritime Safety Agency's (EMSA) CleanSeaNet service, operational since 2007, provides the satellite-based monitoring of the European maritime infrastructure, including the offshore energy infrastructure (the offshore oil and gas platforms, the offshore wind farms, and the subsea pipelines). The CleanSeaNet service processes the Sentinel-1 radar imagery (which can detect oil spills on the sea surface) and the high-resolution optical imagery (which can detect the physical changes to the offshore infrastructure) and provides the results to the member states' coast guards and the offshore infrastructure operators. The service is the exemplar of the satellite-based infrastructure monitoring for the offshore domain, where the in-situ monitoring is difficult and the satellite monitoring is the primary means of the surveillance.

The commercial infrastructure monitoring services, including Maxar's EMI (Energy Monitoring and Intelligence) service and Planet's Planetary Variables service, provide the satellite-based monitoring of the energy infrastructure to the commercial energy companies. Maxar's EMI service uses the WorldView constellation's high-resolution imagery to monitor the transmission corridors, the substations, and the generation facilities, and provides the change detection results to the energy companies' asset management teams. Planet's Planetary Variables service uses the PlanetScope constellation's moderate-resolution imagery (3 metres) to monitor the vegetation conditions along the pipeline rights-of-way and the transmission corridors, providing the vegetation growth alerts to the energy companies' vegetation management teams. These commercial services complement the geospatial office's monitoring, and the office may use the commercial services' products as additional inputs to its own analysis, subject to the license terms and the firewall considerations. The commercial services' strengths are their frequent revisit (Planet's PlanetScope provides daily coverage at 3 metres, enabling the near-real-time monitoring of the vegetation growth) and their analytical platforms (Maxar's EMI provides the change detection results through a web interface, reducing the office's processing burden). The commercial services' limitations are their cost (the services are priced per square kilometre, and the monitoring of a 220-kilometre transmission corridor may cost $15,000 to $25,000 per year), their license terms (the services' license terms may restrict the redistribution of the products to the energy sector entity or the regulator), and their coverage (the services may not cover the remote infrastructure corridors that the office must monitor). The office's decision to use the commercial services is a cost-benefit analysis: the commercial services' frequent revisit and analytical platform may reduce the office's processing burden and improve the monitoring's timeliness, but the cost and the license terms must be weighed against the benefit.

21.9 Worked Example: Detecting Vegetation Encroachment on a Transmission Corridor

The hazard specialist conducts the monthly energy infrastructure monitoring cycle for the jurisdiction's principal transmission corridor: the 220-kilometre, 500-kV transmission line from the Highmoor Power Station to the Portsmith Substation, which serves approximately 320,000 customers in the metropolitan region. The corridor's right-of-way is 60 metres wide (30 metres on each side of the transmission line), and the vegetation clearance zone is 15 metres (the minimum

distance between the vegetation and the transmission line, specified by the sectoral regulator's safety standard).

The specialist retrieves the latest Maxar WorldView-3 image (acquired 3 July 2024, 0.31-metre panchromatic and 1.24-metre multispectral, 5 per cent cloud cover) from the processing engine. The previous month's image (acquired 5 June 2024, same specifications) is retrieved as the reference. Both images have been processed (georeferenced, calibrated to surface reflectance, mosaicked, cloud masked) by the data pipeline operator, and the technical lead has validated the processing outputs.

The specialist configures the change detection plug-in. The change detection algorithm is set to the object-based approach (more robust to the radiometric differences between the two acquisitions, which is important because the two acquisitions are one month apart and the sun angle and the atmospheric conditions have changed). The change threshold is set to 0.15 (the standard parameter for the WorldView-3 multispectral bands, calibrated by the technical lead). The buffer around the transmission line is set to 30 metres (the right-of-way), with a sub-buffer at 15 metres (the clearance zone) for the vegetation encroachment detection.

The specialist executes the change detection run. The run takes approximately 90 minutes (the 220-kilometre corridor at 0.31-metre resolution is a substantial dataset, and the object-based segmentation and the object comparison are computationally intensive). The run identifies 12 changes within the right-of-way: 7 vegetation encroachment changes, 3 third-party construction changes, 1 ground subsidence change, and 1 false alarm (a shadow change caused by the different sun angle between the two acquisitions).

The specialist reviews each detected change. The 7 vegetation encroachment changes are reviewed against the NDVI (computed from the near-infrared and the red bands): all 7 show an NDVI increase of more than 0.15, confirming the vegetation growth. The specialist cross-references the changes with the transmission line's tower locations (from the infrastructure asset registry): 4 of the 7 encroachments are within the 15-metre clearance zone (high-priority), and 3 are within the 30-metre right-of-way but outside the clearance zone (moderate-priority). The 3 third-party construction changes are reviewed against the high-resolution panchromatic image: all 3 show a new structure (a shed, a fence, and a building under construction) within the right-of-way. The specialist cross-references the changes with the cadastre: the shed and the fence are on private parcels adjacent to the right-of-way, and the building is on a parcel that has been subdivided for residential development. The 1 ground subsidence change is reviewed against the DEM: the subsidence is approximately 0.3 metres over a 50-metre section of the corridor, potentially indicating slope instability. The 1 false alarm (the shadow change) is identified and excluded from the alert list.

The specialist computes the vulnerability scores for the 11 confirmed changes (the 7 vegetation encroachments + the 3 third-party constructions + the 1 subsidence). The 4 high-priority vegetation encroachments (within the clearance zone, near critical towers) score 3.2 to 4.0 (high-priority). The 3 moderate-priority vegetation encroachments (outside the clearance zone, near non-critical towers) score 1.0 to 1.4 (low-priority). The 3 third-party constructions score 2.1 to 2.8 (moderate-priority, because the construction is within the right-of-way but not yet within the clearance zone). The 1 subsidence scores 3.5 (high-priority, because the subsidence is near a critical tower and the slope instability could affect the tower's foundation).

The specialist generates the products. The infrastructure risk map (Security Confidential) shows the 220-kilometre transmission corridor, the 12 detected changes (colour-coded by vulnerability score), and the tower locations. The encroachment alerts (Security Confidential) identify the 8 high-priority changes (4 vegetation encroachments + 1 subsidence + 3 third-party constructions) with the supporting imagery (the before-and-after image pairs, annotated to show the encroachment), the vulnerability scores, and the recommended actions (vegetation clearing for the encroachments, geotechnical investigation for the subsidence, legal action for the third-party constructions). The maintenance advisories (Security Confidential) identify the corridor sections at risk from the ground conditions (the subsidence section and the sections with the slope instability potential).

The specialist submits the products to the production operator. The production operator produces the products, applies the Security Confidential classification marking (in the header and footer of every page, in the title block of the map, and in the metadata of the GeoTIFF and the Shapefile), and submits the draft to the office manager. The office manager reviews the products (verifying the classification marking, the analytical content, and the vulnerability scores) and authorizes the release. The products are delivered to the energy sector entity's operations centre through the secure data interface, and the entity's maintenance dispatch uses the 8 high-priority alerts to schedule the field response: the vegetation clearing crews for the 4 vegetation encroachments, the geotechnical team for the subsidence investigation, and the legal team for the third-party construction enforcement.

This worked example illustrates the energy infrastructure domain's full method: the high-resolution imagery acquisition (Maxar WorldView-3, 0.31m), the object-based change detection (threshold 0.15, 90-minute processing), the specialist's review (12 changes, 11 confirmed, 1 false alarm excluded), the cross-referencing with the asset registry (tower locations, criticality ratings), the vulnerability scoring (Criticality Weight × Proximity Factor × Change Type Factor, scores 1.0 to 4.0), and the product generation (risk map, encroachment alerts, maintenance advisories, all Security Confidential). The domain's products provide the energy sector entity with the intelligence it needs to maintain the transmission corridor's safety and reliability, and the Security Confidential classification ensures that the infrastructure's geolocation and vulnerabilities are protected from disclosure. The domain's monthly cadence, the object-based change detection, the vulnerability scoring, and the classification regime together constitute the office's energy infrastructure intelligence capability, and the capability is the basis for the energy sector entity's maintenance decisions, the sectoral regulator's compliance verification, and the emergency management authority's incident response. The 2003 Northeast Blackout, triggered by a transmission line sagging into overgrown vegetation, demonstrated the catastrophic consequences of undetected encroachment: 55 million people without power, $6 billion in economic damage, and a cascading failure that took four days to fully restore. The geospatial office's energy infrastructure intelligence is the mechanism by which such encroachments are detected before they cause failures, and the detection is the domain's contribution to the energy system's reliability and the public's safety, and the contribution is the reason the domain exists.

Chapter 22 — Minerals and Critical Resources Intelligence

22.1 The Domain Defined

Minerals and critical resources intelligence is the production of geological maps, prospectivity assessments, and resource estimates supporting the mineral sector and the strategic stockpile. The domain is the commercial arm of the geospatial office: it generates the revenue that sustains the public-service functions (flood risk, food security, energy infrastructure), and it is the domain where the commercial firewall is most actively exercised. The domain serves the commercial customers (mining and exploration companies that purchase the office's mineral intelligence) and the strategic stockpile authority (the government body that manages the critical minerals stockpile for national security) as its principal consumers.

Critical minerals — the minerals essential to the economy and the national security whose supply is at risk — have become a strategic priority for governments worldwide. The energy transition (from fossil fuels to renewable energy and electric vehicles) requires lithium, cobalt, nickel, graphite, manganese, and rare earth elements in quantities far exceeding the current production. The digital economy requires semiconductors (silicon, germanium, gallium), and the defence industry requires titanium, tungsten, antimony, and beryllium. The United States' List of Critical Minerals (50 minerals, established by the Energy Act of 2020 and updated by the USGS in 2022), the European Union's List of Critical Raw Materials (34 materials, established in 2011 and updated in 2023 by the Critical Raw Materials Act), and the Australian Critical Minerals List (26 minerals, established by the Department of Industry, Science, and Resources) all identify the minerals that the geospatial office's critical minerals inventory must cover.

The domain is conducted on a project-based cadence for the commercial customers, with annual standing products (the national geological map series, the critical minerals inventory). The domain operates under a mixed classification regime: base geology is Public, commercial commissions are Confidential, and the critical minerals inventory is Security Confidential. The mixed regime is the hardest records-management problem in the office, because the three classification levels apply to products that are produced by the same specialist, using the same plug-in, from data that may overlap.

The domain is bounded by three interfaces. The intelligence-regulation boundary separates it from the sectoral regulator's regulatory function (Chapter 26): the geospatial office produces the intelligence commercially, and the regulator regulates the sector. The intelligence-decision division separates it from the strategic stockpile authority's stockpile function: the office produces the inventory, and the authority manages the stockpile. The intelligence-exploration boundary separates it from the commercial minerals firms' exploration function: the office produces the prospectivity intelligence, and the firms conduct the exploration (the drilling, the trenching, the bulk sampling that confirms the mineral deposit's existence and grade).

22.2 The Scope

The minerals and critical resources domain covers three activities: prospectivity assessment, resource estimation, and the critical minerals inventory.

Prospectivity assessment is the assessment of the mineral potential of the geological formations from the geological, geophysical, and geochemical data. The assessment is conducted through the prospectivity plug-in of the modeling workbench (Chapter 10, section 10.5), and the assessment produces a prospectivity map: a raster or a vector dataset that shows, for each location in the study area, the probability that a mineral deposit of the target type occurs at that location. The prospectivity map is the domain's primary product: the commercial customers use the map to prioritize their exploration targets, and the stockpile authority uses the map to assess the jurisdiction's mineral resource potential.

Resource estimation is the estimation of the quantity and quality of the mineral resource from the prospectivity assessment and the geological data. The estimation is conducted through the geostatistical resource estimation models (ordinary kriging, indicator kriging, conditional simulation), and the estimation follows the international reporting standards (the JORC Code in Australia, NI 43-101 in Canada, the SME Standard in the United States), which specify the data quality requirements, the estimation methods, and the reporting categories (inferred, indicated, measured resources; probable and proven reserves). The resource estimation is the domain's analytical product: the commercial customers use the estimate to assess the economic viability of the deposit, and the stockpile authority uses the estimate to assess the jurisdiction's resource base.

The critical minerals inventory is the inventory of the critical minerals resources, aggregated from the resource estimates for the minerals on the government's critical minerals list. The inventory provides the strategic stockpile authority with the information needed to plan the stockpile's composition (which minerals to stockpile, in what quantities, at what rates of acquisition), to assess the supply risk for each mineral (the domestic production vs. the import dependence, the concentration of the global supply in a small number of producer countries, the geopolitical risk of the supply chain), and to identify the minerals where the domestic resource base is insufficient and where the exploration investment should be prioritized.

The scope does not include the mineral exploration (the drilling, the trenching, the bulk sampling that confirms the deposit), which is the commercial minerals firms' function. It does not include the mineral regulation (the mining permits, the environmental compliance, the safety standards), which is the sectoral regulator's function. It does not include the stockpile management (the physical storage, the rotation, the quality control), which is the strategic stockpile authority's function. An office that expands its scope to include the exploration, the regulation, or the stockpile management exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

22.3 The Method

The method of the minerals and critical resources domain is the mineral prospectivity modeling, supplemented by resource estimation and critical minerals inventory aggregation.

The prospectivity modeling combines the geological, geophysical, and geochemical data into a mineral potential map. Two approaches are used. The knowledge-driven approach encodes the geological expert's criteria into a set of weights. The expert identifies the geological features that are associated with the target mineral deposit type — for example, for porphyry copper deposits: intrusive rocks of intermediate composition (granodiorite, quartz monzonite), hydrothermal

alteration zones (potassic, phyllic, argillic), geochemical anomalies in copper and molybdenum, magnetic anomalies associated with the intrusive, and structural features (faults, fractures) that controlled the mineralizing fluids — and the model combines the weighted features into a prospectivity map.

The weights-of-evidence method, developed by Bonham-Carter et al. in 1989, is the principal knowledge-driven approach. The method uses Bayes' theorem to compute the weights from the spatial association between the known deposits and the geological features. For each feature, the positive weight (W+) is the log odds of a deposit occurring given the feature's presence, and the negative weight (W-) is the log odds of a deposit occurring given the feature's absence. The weights are computed from the training data: the known deposits (the positive examples) and a set of randomly selected non-deposit locations (the negative examples). The prospectivity score at each location is the sum of the weights for all features present at that location: Prospectivity = Σ(W+_i) for all features i present at the location. The score is then normalized to a 0-to-1 scale for the prospectivity map.

The fuzzy logic method, developed by An et al. in 1991, uses fuzzy membership functions to represent the uncertainty in the geological expert's criteria. Each feature is assigned a fuzzy membership value between 0 and 1, representing the degree to which the feature is favourable for the deposit. The fuzzy memberships are combined using fuzzy operators: fuzzy AND (the minimum, representing the most conservative combination), fuzzy OR (the maximum, representing the most permissive combination), and fuzzy gamma (a weighted combination of the AND and the OR, allowing the expert to control the balance between the conservative and the permissive).

The data-driven approach uses known deposits to train a machine learning model. The random forest classifier, developed by Breiman in 2001, is an ensemble of decision trees that is robust to overfitting and can handle high-dimensional data with mixed variable types. The random forest is trained on the features (the geological, geophysical, and geochemical data) at the known deposit locations (the positive examples) and at the randomly selected non-deposit locations (the negative examples), and the trained model predicts the probability of a deposit at each location in the study area. The random forest's advantages are its robustness (it does not overfit as easily as a single decision tree), its ability to handle non-linear relationships (the deposit's occurrence may depend on the interaction of multiple features), and its interpretability (the feature importance scores show which features contribute most to the prediction).

The resource estimation is conducted through the geostatistical methods. Ordinary kriging estimates the grade (the concentration of the target mineral) at unsampled locations from the grades at the sampled locations (the drill holes, the trench samples), weighted by the spatial covariance of the grades (the variogram, which describes how the grade's similarity decreases with distance). The kriging estimate is the best linear unbiased estimator: it minimizes the estimation error variance under the constraint that the estimator is unbiased and linear in the data. Indicator kriging estimates the probability that the grade exceeds a specified cutoff, which is useful for the resource classification (the probability that the grade is above the economic cutoff determines whether the resource is classified as inferred, indicated, or measured). Conditional simulation generates multiple equally probable realizations of the grade distribution, allowing the assessment of the uncertainty in the

resource estimate (the range of the possible tonnages and grades, which the commercial customers need for their investment decisions).

The resource estimation follows the international reporting standards. The Joint Ore Reserves Committee (JORC) Code, developed in Australia and first published in 1989, specifies the data quality requirements (the drilling density, the sampling methods, the assay quality), the estimation methods (the geostatistical methods, the validation against the data), and the reporting categories (inferred resource: low confidence, based on limited data; indicated resource: moderate confidence, based on sufficient data; measured resource: high confidence, based on detailed data; probable reserve: the economically mineable part of an indicated or measured resource; proven reserve: the economically mineable part of a measured resource). The Canadian National Instrument 43-101 (NI 43-101), established in 2001, and the Society for Mining, Metallurgy, and Exploration (SME) Standard, used in the United States, follow similar principles. The commercial customers require the resource estimates to comply with these standards, because the estimates are used in the regulatory filings (the mineral resource and reserve statements that the publicly listed mining companies must file with the securities regulators) and in the investment decisions (the bankable feasibility studies that the mining companies' investors require).

The validation is conducted against the known deposits recorded in the geological survey's mineral occurrence database. The validation compares the model's prospectivity predictions with the locations of the known deposits: the prospectivity accuracy (the percentage of the known deposits that fall within the model's high-prospectivity areas) and the prediction ratio (the ratio of the known deposit density in the high-prospectivity areas to the known deposit density in the overall study area) are the accuracy metrics, with the tolerances specified in Chapter 10, section 10.6 (70 per cent prospectivity accuracy, 3.0 prediction ratio).

22.4 The Products

The minerals and critical resources domain produces four product types, each serving a different consumer and a different purpose.

The prospectivity maps are the domain's standing product, updated annually. The maps show the mineral potential of the geological formations, using the warm-cool gradient (dark red for the highest prospectivity, orange for high, yellow for moderate, light blue for low, dark blue for the lowest), with the geological formations overlaid in muted browns and greens and the known deposits marked as stars. The maps are produced for the mineral types on the government's critical minerals list and for the mineral types that the commercial customers commission. The standing prospectivity maps (for the critical minerals) carry the Public classification (for the base geology) or the Security Confidential classification (for the critical minerals inventory). The commissioned prospectivity maps (for the commercial customers) carry the Confidential classification.

The resource estimates are the domain's analytical product, produced for the commercial customers and the stockpile authority. The estimates follow the international reporting standards (JORC, NI 43-101, SME) and specify the resource category (inferred, indicated, measured), the tonnage (the estimated quantity of the mineralized material, in million tonnes), the grade (the estimated concentration of the target mineral, in per cent or in parts per million), and the contained metal (the tonnage multiplied by the grade, in tonnes of metal). The resource estimates are produced

from the drill hole data, the trench sample data, and the bulk sample data provided by the commercial customers or by the geological survey's own exploration programmes.

The critical minerals inventory is the domain's strategic product, produced annually for the strategic stockpile authority. The inventory aggregates the resource estimates for the minerals on the government's critical minerals list, providing the authority with the domestic resource base for each critical mineral, the production rate (the annual production from the operating mines), the import dependence (the share of the consumption that is met by imports), and the supply risk assessment (the concentration of the global supply in a small number of producer countries, the geopolitical risk of the supply chain, and the potential for the supply disruption). The inventory carries the Security Confidential classification, because the disclosure of the jurisdiction's critical minerals vulnerabilities could inform strategic adversaries.

The commercial commission reports are the domain's commercial product, produced on a project basis for the commercial customers. The reports comprise the prospectivity map (the warm-cool gradient with the geological formations and the known deposits), the resource estimate (the tonnage, the grade, and the contained metal, following the JORC or NI 43-101 reporting standard), the geological interpretation (the explanation of the deposit's geological setting, the mineralizing processes, and the exploration potential), and the recommendations (the suggested exploration targets, the suggested drilling programmes, and the suggested sampling strategies). The reports carry the Confidential classification, with the customer identity and the product identifier in the marking. The reports are the product that generates the revenue that sustains the public-service functions, and the reports are the product that the commercial firewall is designed to govern.

22.5 The Consumers

The minerals and critical resources domain serves three consumers, each with different needs and different commercial or strategic contexts.

The commercial customers are the principal consumers. The customers are the mining and exploration companies that purchase the office's mineral intelligence to support their exploration decisions. The exploration decision is: where to drill? Drilling is expensive (approximately $150 to $500 per metre for diamond drilling, depending on the location and the depth), and a typical exploration programme may cost $5 million to $20 million. The prospectivity map reduces the exploration risk by identifying the areas with the highest probability of a deposit, allowing the customer to focus the drilling on the most promising targets. The commercial customers require the intelligence in formats that support their exploration planning (GeoTIFF for their GIS systems, PDF for their technical reports, and the JORC/NI 43-101-compliant resource estimate for their regulatory filings and their investor presentations).

The strategic stockpile authority is a secondary consumer. The authority uses the critical minerals inventory for the stockpile management: the inventory's domestic resource base assessment informs the authority's decision on which minerals to stockpile domestically (the minerals with a sufficient domestic resource base can be stockpiled from the domestic production) and which minerals to stockpile through the international procurement (the minerals with an insufficient domestic resource base must be procured from the international market, subject to the supply risk). The authority requires the inventory in a format that supports the stockpile planning (the resource

estimates by mineral, by region, and by deposit type, with the confidence levels and the supply risk assessments).

The public is a tertiary consumer. The public uses the public base geology maps for the geological awareness: the geological maps published by the geological survey are used by the academic community (for the geological research), by the planning community (for the land-use planning, the groundwater assessment, and the hazard assessment), and by the general public (for the educational and the recreational purposes, such as the rockhounding and the fossil collecting). The public base geology maps carry the Public classification and are available through the geological survey's data portal.

22.6 The Cadence

The minerals and critical resources domain operates on a project-based cadence for the commercial customers, with annual standing products. The two cadences run in parallel, and the specialist manages both.

The project-based cadence comprises the commercial commission's lifecycle. The lifecycle begins with the commission request (the commercial customer contacts the office manager to request a prospectivity assessment for a specific area and a specific mineral type). The office manager evaluates the request (the area, the mineral type, the data availability, the delivery timeline, the price), accepts or declines the commission, and opens the commercial charge code. The specialist conducts the data integration (the geological, geophysical, and geochemical data are acquired and processed), the model run (the prospectivity model is configured, executed, and validated), and the product generation (the prospectivity map, the resource estimate, and the commission report are produced). The office manager reviews the products under the commercial release authorization procedure (section 13.4), and the production operator releases the products through the commercial channel. The typical project duration is 4 to 12 weeks, depending on the area's size, the data's availability, and the model's complexity.

The annual cadence comprises the standing products' production cycle. The cycle begins with the data update (the geological survey's new mapping, the new geophysical surveys, and the new exploration data are integrated into the prospectivity model), the model recalibration (the model's weights or the model's parameters are recalibrated using the new known deposits discovered during the year), the validation (the recalibrated model is validated against the updated known deposit database), and the product generation (the updated prospectivity maps, the updated resource estimates, and the updated critical minerals inventory are produced). The annual cycle is conducted in the first quarter of each calendar year, and the products are delivered to the stockpile authority by the end of the first quarter.

The project-based cadence is maintained by the resource economics specialist, who produces the commercial commission reports, and the office manager, who authorizes the commercial releases. The annual cadence is maintained by the resource economics specialist, who produces the standing products, and the technical lead, who validates the products. The two cadences' coordination is the specialist's responsibility: the specialist must balance the commercial commissions' deadlines (which generate the revenue) with the standing products' schedule (which serves the stockpile authority), and

the balance is the operational discipline that ensures both cadences are maintained without one compromising the other.

22.7 The Classification Regime

The minerals and critical resources domain operates under a mixed classification regime: the base geology is Public, the commercial commissions are Confidential, and the critical minerals inventory is Security Confidential. The mixed regime is the hardest records-management problem in the office because the three classification levels apply to products that are produced by the same specialist, using the same plug-in, from data that may overlap.

The base geology is Public. The public geological maps — the formation boundaries, the rock types, the structural features (faults, folds), the stratigraphic columns — are the foundation of the public's geological literacy and the academic community's geological research. The geological surveys (the USGS, Geoscience Australia, the BGS) publish their geological maps freely, and the geospatial office's public base geology maps are consistent with this practice. The Public classification ensures that the geological information is available to the academic community, the planning community, and the general public, and the availability is the condition of the public's engagement with the mineral sector.

The commercial commissions are Confidential. The commercial customer's identity, the specific area of interest, the mineral type, and the model's findings are commercially sensitive, because the disclosure could compromise the customer's competitive position. If a competitor learns that Meridian Resources has commissioned a prospectivity assessment for the northern highlands, the competitor may infer that Meridian has identified a promising target and may race to acquire the adjacent exploration licenses. The Confidential classification, with the customer identity and the product identifier in the marking, restricts the commission report to the customer and the office, and the restriction is the mechanism that protects the customer's investment in the intelligence.

The critical minerals inventory is Security Confidential. The inventory identifies the jurisdiction's domestic resource base for each critical mineral, the production rate, the import dependence, and the supply risk. The disclosure of the inventory could inform strategic adversaries about the jurisdiction's resource vulnerabilities: an adversary that learns that the jurisdiction is 90 per cent import-dependent on a specific critical mineral may target the supply chain (through export restrictions, through the coercion of the producer countries, or through the sabotage of the transportation routes) to disrupt the jurisdiction's economy or defence industry. The Security Confidential classification restricts the inventory to the security-cleared personnel (the stockpile authority's security-cleared staff, the defence ministry's logistics planners, and the intelligence agencies).

The mixed regime is the reason the commercial release procedure (section 13.4) exists as a distinct procedure. The procedure ensures that the commercial products are released only through the commercial channel, only to the authorized customer, and only after the office manager's authorization. The procedure also ensures that the commercial products are not commingled with the public base geology or the Security Confidential critical minerals inventory in the archive or in the distribution system. The production operator's classification marking verification (section 16.4) checks every product for the correct marking, and the office manager's review verifies the marking

before release. The mixed regime is non-negotiable, and a product that is misclassified has compromised either the public awareness, or the customer's competitive position, or the jurisdiction's strategic security.

22.8 The Comparative Practice

The United States Geological Survey (USGS), established in 1879, is the world's pre-eminent geological survey and the exemplar of the minerals-funded public geoscience model. The USGS Mineral Resources Program, through its National Minerals Information Center, compiles and publishes the mineral commodity summaries for approximately 90 mineral commodities, covering the production, the consumption, the imports, the exports, and the stocks for the United States and the world. The USGS also conducts the mineral resource assessments for the United States, using the three-part assessment methodology developed by the Survey (the geological analysis, the grade-tonnage models, and the spatial analysis), and publishes the assessments through the Scientific Investigations Map series. The USGS Earth Resources Observation and Science (EROS) Center sells value-added products and custom services on a cost-recovery basis, and the revenue is deposited in the Survey's revolving fund (authorized by 43 U.S.C. § 364b). The Survey's Mineral Resources Data System (MRDS) is the world's largest mineral occurrence database, with records for over 260,000 mineral occurrences in the United States and worldwide.

Geoscience Australia, established in 2001 through the merger of AGSO and AUSLIG, is the exemplar of the minerals-commercial model in a jurisdiction with a substantial mineral sector. Australia is the world's largest producer of iron ore, lithium, and bauxite, and the second-largest producer of gold, and the mineral sector contributes approximately 10 per cent of Australia's GDP. Geoscience Australia's Minerals Systems program conducts the mineral prospectivity assessments for the Australian continent, using the weights-of-evidence and the machine learning methods, and publishes the assessments through the Australian Mineral Exploration Reviews and the mineral potential maps. The agency's Exploring for the Future programme (launched in 2016, with a $100.5 million budget) conducted the regional-scale geological, geophysical, and geochemical surveys of the underexplored regions of northern Australia, and the data is publicly available through the agency's data portal. The agency sells mineral intelligence through its minerals division, and the revenue supports the agency's public geoscience programmes.

The British Geological Survey (BGS), established in 1835 as the Geological Survey of Great Britain (the world's oldest national geological survey), is the exemplar of the commercial cost-recovery model. BGS operates as a research centre of the Natural Environment Research Council (NERC), and its commercial division provides custom geological reports, site investigations, and geotechnical assessments to the construction, mining, and energy sectors. The BGS publishes the DiGMapGB (the Digital Geological Map of Great Britain) at 1:50,000 and 1:10,000 scales, the mineral resource maps for the United Kingdom, and the mineral planning factsheets for the principal minerals. The BGS's commercial revenue supports the Survey's public science, including the national groundwater monitoring network and the British National Seismological Database.

These three surveys confirm the domain's doctrine. The prospectivity assessment is the foundation: the USGS's three-part assessment, Geoscience Australia's weights-of-evidence and machine learning models, and the BGS's mineral resource maps all begin with the integration of the

geological, geophysical, and geochemical data into a mineral potential assessment. The resource estimation is the analytical core: all three surveys follow the international reporting standards (JORC in Australia, NI 43-101 in Canada, SME in the United States) for the resource estimates. The commercial commission is the revenue-neutral value: the USGS's cost-recovery sales through EROS, Geoscience Australia's minerals division, and the BGS's commercial division all generate the revenue that supports the public geoscience. The three surveys also confirm the mixed classification regime: the public geological maps (the USGS's Scientific Investigations Maps, Geoscience Australia's mineral potential maps, the BGS's DiGMapGB) are public, the commercial products are sold under commercial terms, and the critical minerals information is restricted.

22.9 Worked Example: A Commercial Prospectivity Assessment for a Lithium Commission

A mining company, LithiumQuest Resources, has commissioned a prospectivity assessment for hard rock lithium (spodumene-bearing pegmatites) in a 2,000-square-kilometre area of the eastern highlands. The commission is for a custom prospectivity map, a resource estimate, and a commission report, with a delivery timeline of 10 weeks and an agreed price of $65,000. The office manager accepts the commission, opens the commercial charge code COMM-2024-023, and assigns the resource economics specialist to the project.

The specialist begins with the data integration. The geological data includes the 1:50,000 geological map (from the geological survey, showing the formation boundaries, the rock types, and the structural features), the geochemical data (from the geological survey's stream sediment survey, with the lithium, cesium, and tantalum concentrations for approximately 2,500 sample sites), and the geophysical data (the aeromagnetic survey, from the geological survey's geophysical archive, flown at 250-metre line spacing and 60-metre terrain clearance, providing the total magnetic intensity at 50-metre grid resolution). The hyperspectral data (from the ASTER satellite, with the 14 spectral bands in the visible, near-infrared, and shortwave-infrared, used to map the alteration minerals associated with the lithium-cesium-tantalum pegmatites) is acquired from the NASA Earthdata portal. All data is acquired under the commercial charge code COMM-2024-023, and the license terms are recorded in the catalog and archive submodule.

The specialist configures the prospectivity model. The target deposit type is the lithium-cesium-tantalum (LCT) pegmatite, which is associated with peraluminous granites (S-type granites), and the expert criteria include: the proximity to the peraluminous granite intrusions (within 5 kilometres), the presence of the lithium-cesium-tantalum geochemical anomaly (the stream sediment samples with lithium > 200 ppm, cesium > 20 ppm, or tantalum > 10 ppm), the presence of the structural features (the faults and the shear zones that controlled the pegmatite emplacement), and the presence of the alteration minerals (the muscovite and the kaolinite, detectable in the ASTER shortwave-infrared bands). The specialist uses the weights-of-evidence method (Chapter 10, section 10.5), with the training data from the geological survey's mineral occurrence database (23 known LCT pegmatite occurrences in the study area).

The specialist executes the model run. The weights-of-evidence computation produces the W+ and W- weights for each feature: the peraluminous granite proximity (W+ = 1.8, W- = -0.3), the lithium geochemical anomaly (W+ = 2.1, W- = -0.4), the cesium anomaly (W+ = 1.5, W- = -0.2), the

tantalum anomaly (W+ = 1.2, W- = -0.1), the structural features (W+ = 0.9, W- = -0.2), and the alteration minerals (W+ = 1.3, W- = -0.2). The prospectivity score at each location is the sum of the weights for all features present at that location, normalized to a 0-to-1 scale. The model run takes approximately 4 hours (the 2,000-square-kilometre area at 50-metre grid resolution is a substantial dataset, and the weights-of-evidence computation includes the spatial association analysis for each feature).

The specialist validates the model run against the known deposits. The prospectivity accuracy is 82 per cent (19 of the 23 known LCT pegmatite occurrences fall within the high-prospectivity areas, above the 70 per cent tolerance). The prediction ratio is 5.1 (the known deposit density in the high-prospectivity areas is 5.1 times the known deposit density in the overall study area, above the 3.0 tolerance). The specialist documents the validation results in the provenance record, under the commercial charge code COMM-2024-023.

The specialist produces the prospectivity map. The map uses the warm-cool gradient (dark red for the highest prospectivity, 0.8 to 1.0; orange for high, 0.6 to 0.8; yellow for moderate, 0.4 to 0.6; light blue for low, 0.2 to 0.4; dark blue for the lowest, 0.0 to 0.2), with the geological formations overlaid in muted browns and greens at 30 per cent transparency, the known LCT pegmatite occurrences marked as black stars (10-point), and the peraluminous granite intrusions outlined in black (2-point). The map identifies three high-priority targets (areas with the prospectivity score above 0.8 that do not contain a known deposit): Target A (12 square kilometres, in the northern section, with all six features present), Target B (8 square kilometres, in the central section, with five of six features present), and Target C (5 square kilometres, in the southern section, with five of six features present).

The specialist produces the commission report. The report comprises the prospectivity map, the methodology section (the weights-of-evidence method, the training data, the W+ and W- weights for each feature, the validation results), the resource estimate (the inferred resource for each target, based on the expected tonnage and grade for the LCT pegmatite deposit type, following the JORC Code's inferred resource category), and the recommendations (the suggested drilling programme for each target: the diamond drilling at 100-metre spacing along the pegmatite's strike, to a depth of 200 metres, with the estimated cost of $2.5 million per target).

The specialist applies the Confidential classification marking, with the customer identity ("LithiumQuest Resources") and the product identifier ("COMM-2024-023") in the marking. The specialist submits the commission report to the production operator. The production operator produces the report, and the office manager reviews the report under the commercial release authorization procedure (section 13.4): the firewall compliance (the product was produced from commercial data sources and commercial model runs, charge code COMM-2024-023), the classification marking (Confidential with the customer identity and product identifier), and the commercial release log is signed. The production operator releases the report to LithiumQuest Resources through the commercial distribution channel.

This worked example illustrates the minerals and critical resources domain's full method: the data integration (the geological, geochemical, geophysical, and hyperspectral data), the prospectivity modeling (the weights-of-evidence method with the LCT pegmatite expert criteria, the W+ and W- weights for each feature, the 4-hour model run), the validation (82 per cent accuracy above the 70 per

cent tolerance, 5.1 prediction ratio above the 3.0 tolerance), the product generation (the prospectivity map with the three high-priority targets, the resource estimate following the JORC Code, the commission report with the drilling recommendations), and the classification regime (the Confidential marking with the customer identity, the commercial charge code, the commercial release authorization, and the commercial channel release). The domain's products provide LithiumQuest Resources with the intelligence it needs to prioritize its exploration drilling, and the intelligence is the commercial value that generates the revenue sustaining the public-service functions. The domain's contribution to the revenue-neutral model is the reason the domain exists, and the commercial firewall's maintenance is the condition of the domain's credibility.

Chapter 23 — The Cadastral Authority Interface

Part V — Sectoral Interfaces: The inter-agency boundaries with the geospatial office's principal partners, each stated as what the office does, what the partner does, and where the boundary lies.

23.1 The Interface Defined

The cadastral authority interface is the boundary between the office that records land and the office that analyses it. The cadastral authority maintains the legally binding map of land parcels, their boundaries, their attributes, and their ownership, and it does so under statutory authority that makes the cadastre the definitive record for property rights, transactions, taxation, and dispute resolution. The geospatial office consumes that record as a base layer and produces analytical intelligence about the land the cadastre describes: flood inundation maps joined to parcel identifiers, crop condition assessments joined to field boundaries, infrastructure risk assessments joined to property addresses, mineral prospectivity studies joined to exploration licence areas. The two offices are operationally interdependent because the cadastre is the base layer for every analytical product, and they are institutionally distinct because the legal authority to record and the analytical independence to assess are different integrity requirements that should not be combined in a single office.

The boundary is a functional boundary, not a bureaucratic one. The cadastral authority's function requires legal authority, statutory process, and defensibility in court, because the cadastre is the record on which property rights, transactions, and disputes are decided. The authority's surveyors are licensed under statute, the boundary determinations follow statutory procedures (the traverse, the monument, the deed, the title), and the records are admissible in court as evidence of the legal boundary. The geospatial office's function requires analytical independence, methodological transparency, and the freedom to produce intelligence that may contradict the assumptions embedded in the cadastre. If the office's flood model reveals that a parcel registered as upland is in fact within the floodplain, the office must be free to publish that finding without seeking the cadastral authority's permission, and the finding must be intelligible to the consumer as an analytical product, not as a cadastral determination. The two functions have different integrity requirements, and combining them in a single office creates pressure either to make the cadastre serve the analysis (the cadastral record is adjusted to fit the model's output) or to make the analysis serve the cadastre (the model's output is suppressed or modified to fit the cadastral record). The boundary prevents both pressures.

The boundary is the boundary between recording and analysing, and it is restated in this chapter because every chapter that follows in Part V (the sectoral interfaces with the emergency management authority, the agriculture department, the energy sector regulator, the strategic reserve authority, and the commercial customers) depends on the cadastre as a base layer. The flood inundation map is meaningful when it can be associated with specific parcels; the crop condition assessment is meaningful when it can be associated with specific fields; the infrastructure risk assessment is meaningful when it can be associated with specific properties; the mineral prospectivity study is meaningful when it can be associated with specific exploration licences. The cadastre feed is the mechanism by which the geospatial office's analytical products are joined to the legal and administrative reality of the land, and the interface with the cadastral authority is the mechanism by which the feed is maintained.

23.2 The Scope

The cadastral authority interface covers four activities: the cadastre feed's ingestion, the synchronization audit, the discrepancy reporting, and the inter-agency data-sharing framework's maintenance. Each activity is the geospatial office's responsibility, and each is conducted within the boundary the interface defines.

The cadastre feed's ingestion is the operational mechanism of the interface. The feed provides the current state of the cadastre — the parcels, their boundaries, their legal attributes (the parcel identifier, the area, the address, the land use code, the encumbrances, the easements), and the ownership information (the title number, the registered proprietor, the tenure type) — in a standardized format that the geospatial office ingests into its spatial data infrastructure. The feed is maintained by the cadastral authority, updated as land transactions occur (the subdivisions, the consolidations, the boundary adjustments, the transfers of title), and distributed to the geospatial office through a data interface that complies with the national spatial data infrastructure's standards. In jurisdictions that have implemented the Open Geospatial Consortium's Web Feature Service (WFS) for cadastre distribution, the feed is consumed as a streaming service; in jurisdictions that distribute the cadastre as a file geodatabase replica, the feed is consumed as a nightly extract. The OGC WFS 2.0 specification, ISO 19152 Land Administration Domain Model (LADM), and the INSPIRE Annex II Cadastral Parcels theme are the principal international standards that govern the feed's structure, and the geospatial office's ingestion pipeline must implement the relevant standard for its jurisdiction.

The synchronization audit is the operational check that the office's ingested cadastre is consistent with the cadastral authority's current record. The audit is conducted weekly by the data pipeline operator and reviewed by the technical lead, following the procedure specified in Chapter 8, section 8.4, and the job breakdown sheet in Chapter 15, section 15.3. The audit compares the office's cadastre layer against the cadastral authority's published change log, identifies the parcels that have been modified since the last synchronization, and updates the office's layer to reflect the changes. The audit also checks for the structural integrity of the cadastre layer: the topology (no overlapping parcels, no gaps between adjacent parcels), the attribution (every parcel has a valid identifier, area, and land use code), and the projection (the layer is in the office's working projection, with no offset or distortion introduced during the ingestion). The synchronization audit is non-negotiable, because the products produced against an outdated cadastre are wrong in ways that are not visible to the consumer: the flood inundation map joined to an outdated cadastre may show the wrong parcel as flooded, the crop condition assessment joined to an outdated cadastre may attribute the wrong field to the wrong farmer, and the infrastructure risk assessment joined to an outdated cadastre may address the maintenance advisory to the wrong owner.

The discrepancy reporting is the mechanism by which the geospatial office reports to the cadastral authority the discrepancies it discovers between the cadastre and the ground truth. The discrepancies are discovered through the office's analytical work: the flood model may reveal a parcel boundary that does not match the observed terrain (the river has migrated since the boundary was last surveyed, leaving the registered boundary in the channel); the crop condition assessment may reveal a parcel that is registered as agricultural but is being used for another purpose (the field has been converted to a solar farm, but the cadastre still records it as arable); the infrastructure monitoring may

reveal a parcel boundary that is inconsistent with the built infrastructure (a substation enclosure extends beyond the registered parcel boundary, suggesting an encroachment or an unregistered boundary adjustment). The discrepancies are reported to the cadastral authority through the formal discrepancy reporting channel, and the reporting is the mechanism by which the geospatial office contributes to the cadastre's maintenance without assuming the maintenance function.

The inter-agency data-sharing framework is the institutional arrangement that governs the interface. The framework is typically a memorandum of understanding or a service-level agreement between the geospatial office and the cadastral authority, established under the national spatial data infrastructure's inter-agency coordination mechanism (the Federal Geographic Data Committee in the United States, the INSPIRE national contact point in European Union member states, the Australia and New Zealand Land Information Council, the Singapore Geospatial Collaborative Environment). The framework specifies the cadastre feed's format and update frequency, the discrepancy reporting's channel and turnaround time, the synchronization audit's procedures, and the escalation path for unresolved discrepancies. The framework is the institutional expression of the interface, and it is the framework that governs the relationship across changes of leadership, changes of budget, and changes of political priority. The interface that depends on a personal relationship between the two offices' leaders will erode when the relationship changes; the interface that depends on an institutional framework will survive the changes, because the framework is the framework that the two offices' mandates require them to maintain.

The scope does not include surveying, which is the cadastral authority's function and requires statutory authority and licensed surveyors. The office does not survey parcels, does not determine boundaries, does not monument corners, and does not produce survey plans. The scope does not include the land registry, which is the cadastral authority's (or, in separated-authority jurisdictions, the land registry's) function and requires statutory authority. The office does not register transactions, does not issue titles, and does not maintain the register of proprietors. The scope does not include boundary dispute resolution, which is a legal process conducted in the courts or in a statutory tribunal. The office does not adjudicate disputes, does not provide expert testimony on the legal boundary (although the office's specialists may be called as factual witnesses on the analytical products), and does not mediate between disputing parties. The scope's exclusion of surveying, registration, and dispute resolution is the safeguard of the office's analytical independence. The office that conducts its analytical work without assuming the cadastral functions preserves the separation that allows the analytical findings to be received by the cadastral authority as an independent contribution, not as a self-serving adjustment of the record.

23.3 The Method

The method of the cadastral authority interface is the cadastre feed's ingestion, supplemented by the synchronization audit, the discrepancy reporting, and the framework's maintenance. The method's technical detail is the detail of the cadastre's data model, the ingestion pipeline's design, and the synchronization audit's procedures.

The cadastre's data model is the Land Administration Domain Model (LADM), specified in ISO 19152:2012 and revised in ISO 19152:2024. The model organizes the cadastre into four packages: the Party package (the persons and organisations that hold rights to land), the Administrative package

(the rights, the responsibilities, and the restrictions that apply to land), the Spatial Unit package (the parcels, the buildings, and the spatial units that the rights apply to), and the Source package (the documents — the deeds, the survey plans, the mortgages — that establish the rights). The model is implementation-agnostic: a jurisdiction may implement the model as a 2D vector dataset (the parcels as polygons with the administrative attributes as the columns), as a 3D dataset (the parcels as 3D volumes for the stratum titles and the air-rights parcels), or as a topological dataset (the parcels as faces in a planar partition, with the shared boundaries stored once). The geospatial office's ingestion pipeline must implement the LADM's spatial unit package at a minimum, because the parcel polygon and the parcel identifier are the join keys for every analytical product. The office's ingestion pipeline typically implements the administrative package's basic attributes (the land use code, the area, the address) for the analytical work that depends on those attributes, and the party package's basic attributes (the registered proprietor, the tenure type) for the consumer-specific products that require owner identification.

The ingestion pipeline's design is the design of the data flow from the cadastral authority's distribution interface to the office's spatial data infrastructure. The pipeline comprises four stages: the acquisition (the cadastre feed is fetched from the distribution interface, whether as a WFS GetFeature request, a REST API call, or a file download from an FTP server), the validation (the fetched data is validated against the LADM schema and the office's working projection; the validation checks the geometry's validity, the attribution's completeness, and the projection's correctness), the transformation (the validated data is transformed from the cadastral authority's distribution format to the office's working format — typically a PostGIS or a file geodatabase table — with the projection transformed to the office's working projection and the attributes mapped to the office's schema), and the loading (the transformed data is loaded into the office's spatial data infrastructure, with the previous version archived for the provenance chain). The pipeline is automated through the data pipeline operator's orchestration tool (the FME Server, the Apache Airflow, the GeoServer's GeoWebCache, or a custom Python script using the GDAL/OGR library), and the pipeline's execution is logged for the provenance record.

The synchronization audit's procedures are the procedures specified in Chapter 8, section 8.4, and the job breakdown sheet in Chapter 15, section 15.3. The audit compares the office's cadastre layer against the cadastral authority's published change log for the audit period (typically the previous week), identifies the parcels that have been modified, and updates the office's layer. The change log is the cadastral authority's published list of the parcel modifications, distributed as a WFS transaction log, a REST API query result, or a file extract. The audit's procedures include the change detection (the parcels added, modified, or deleted since the last synchronization), the change application (the office's layer is updated to reflect the changes), the topology check (the updated layer is checked for the overlapping parcels and the gaps), and the attribution check (the updated layer is checked for the missing or invalid attributes). The audit's results are documented in the synchronization audit log, which records the date of the audit, the number of the changes applied, the topology and attribution check results, and the technical lead's authorization. The audit is the operational check that prevents the office from producing intelligence against an outdated cadastre, and the audit's documentation is the audit trail that demonstrates the office's due diligence to the consumers and to the oversight bodies.

The discrepancy reporting's procedures are the procedures by which the geospatial office reports the discrepancies to the cadastral authority. The discrepancies are identified through the office's analytical work, and they are documented in a discrepancy report that comprises the parcel identifier, the discrepancy's description (the boundary mismatch, the missing parcel, the attribution error, the encroachment), the evidence (the analytical product that revealed the discrepancy, with the imagery, the model output, or the field observation that supports the finding), the recommended action (the boundary resurvey, the attribution correction, the encroachment investigation), and the office's contact for the follow-up. The report is submitted to the cadastral authority through the formal discrepancy reporting channel specified in the inter-agency data-sharing framework, and the submission is logged in the discrepancy reporting register. The cadastral authority investigates the discrepancy, resolves it through the statutory process (the resurvey, the administrative correction, the court order), and notifies the geospatial office of the resolution. The office updates its cadastre layer to reflect the resolution, and the resolution is documented in the synchronization audit log.

The framework's maintenance is the ongoing coordination between the geospatial office and the cadastral authority under the inter-agency data-sharing framework. The maintenance comprises the regular liaison meetings (typically quarterly, between the office manager and the cadastral authority's counterpart), the framework's review (typically annual, to update the cadastre feed's format or the discrepancy reporting's procedures as the standards and the technologies evolve), and the escalation handling (the unresolved discrepancies, the data-sharing disputes, and the framework's interpretation disagreements are escalated to the national spatial data infrastructure's coordination body for resolution). The maintenance is the office manager's responsibility, and the maintenance is the institutional work that ensures the interface is maintained as an institutional arrangement, not as a personal relationship.

23.4 The Products

The cadastral authority interface produces no products of its own. The interface's products are the cadastre-dependent analytical products of the other four domains, each joined to the cadastre through the parcel identifier. The interface's contribution to those products is the cadastre base layer that the products are joined to, and the interface's integrity is the integrity of the join.

The flood inundation map (Chapter 19, section 19.4) is the first cadastre-dependent product. The map shows the inundation extent and depth for a specified flood event, with the parcels overlaid and colour-coded by the inundation depth at the parcel's centroid. The map's consumer — the emergency management authority — uses the map to identify the parcels affected by the flood, the parcels' registered proprietors (for the evacuation notifications and the damage assessments), and the parcels' land use codes (for the priority routing: the residential parcels before the agricultural, the agricultural before the industrial). The map's join to the cadastre is the parcel identifier, and the join's integrity is the integrity of the synchronization audit: a parcel that has been subdivided since the last synchronization may be shown as a single large parcel when it is in fact two smaller parcels, and the evacuation notification may be sent to the wrong proprietor.

The crop condition assessment (Chapter 20, section 20.4) is the second cadastre-dependent product. The assessment shows the vegetation condition for each agricultural parcel, derived from the satellite imagery's vegetation indices (NDVI, EVI, LAI) and joined to the cadastre through the parcel

identifier. The assessment's consumer — the strategic reserve authority — uses the assessment to identify the parcels with the stressed crops, the parcels' registered proprietors (for the procurement advisories and the relief programmes), and the parcels' areas (for the production forecasts, which are computed as the area multiplied by the expected yield per hectare). The assessment's join to the cadastre is the parcel identifier, and the join's integrity is the integrity of the synchronization audit: a parcel that has been consolidated with the adjacent parcel since the last synchronization may be shown as two separate parcels when it is in fact one, and the production forecast may double-count the area.

The energy infrastructure risk assessment (Chapter 21, section 21.4) is the third cadastre-dependent product. The assessment shows the infrastructure assets (the substations, the pipelines, the transmission towers) with the risk classification (high, medium, low), joined to the cadastre through the parcel identifier for the asset's location. The assessment's consumer — the energy sector regulator and the infrastructure operator — uses the assessment to identify the parcels with the high-risk assets, the parcels' registered proprietors (for the maintenance access agreements and the easement enforcement), and the parcels' encumbrances (for the easements that grant the operator the access rights and the restrictions that limit the operator's activities). The assessment's join to the cadastre is the parcel identifier, and the join's integrity is the integrity of the synchronization audit: an asset that has been moved to a new parcel since the last synchronization may be associated with the wrong parcel, and the maintenance access agreement may be sought from the wrong proprietor.

The mineral prospectivity study (Chapter 22, section 22.4) is the fourth cadastre-dependent product. The study shows the mineral prospectivity of the geological formations, joined to the cadastre through the parcel identifier for the prospectivity's location and through the exploration licence layer (maintained by the sectoral regulator, not the cadastral authority) for the licence area. The study's consumer — the commercial minerals customer and the strategic stockpile authority — uses the study to identify the parcels with the high prospectivity, the parcels' registered proprietors (for the access negotiations and the surface rights agreements), and the parcels' land use codes (for the land-use constraints on the exploration: the residential and the conservation parcels are excluded from the exploration target). The study's join to the cadastre is the parcel identifier, and the join's integrity is the integrity of the synchronization audit: a parcel that has been reclassified from agricultural to conservation since the last synchronization may be shown as an exploration target when it is in fact excluded, and the customer may invest in the access negotiations for a parcel that cannot be explored.

The cadastre-dependent products' integrity is the interface's integrity, and the interface's integrity is the synchronization audit's integrity. The audit is the operational check that ensures the cadastre layer is current, and the audit's documentation is the audit trail that demonstrates the office's due diligence. An office that produces cadastre-dependent products without conducting the synchronization audit has compromised the products' integrity, and the compromise is not visible to the consumer: the consumer sees the parcel identifier and the registered proprietor on the product, and assumes that the join is correct, when the join may be against an outdated cadastre. The synchronization audit is the procedure that prevents the compromise, and the audit's conduct is the office's responsibility, not the cadastral authority's.

23.5 The Consumers

The cadastral authority interface has no consumers of its own. The interface's consumers are the consumers of the cadastre-dependent analytical products, and each consumer's relationship with the interface is mediated by the product the consumer receives. The interface's integrity is the integrity that the consumers rely on, and the interface's failure is the failure that the consumers experience as a failure of the analytical product.

The emergency management authority is the first indirect consumer. The authority's flood response depends on the flood inundation map (Chapter 19), and the map's join to the cadastre is the join that allows the authority to issue the evacuation notifications to the registered proprietors, to route the rescue resources to the affected parcels, and to assess the damage for the relief programmes. The authority's reliance on the cadastre join is total: a flood inundation map without the parcel join is a hydrological product, not an emergency management product, and the authority cannot use it for the operational response. The authority's interface with the geospatial office is treated in Chapter 24 (the emergency management authority interface), and the cadastral interface is the underlying interface that makes the emergency management interface's products actionable.

The strategic reserve authority is the second indirect consumer. The authority's crop procurement decisions depend on the crop condition assessment (Chapter 20), and the assessment's join to the cadastre is the join that allows the authority to identify the parcels with the stressed crops, the parcels' proprietors (for the procurement advisories), and the parcels' areas (for the production forecasts). The authority's interface with the geospatial office is treated in Chapter 25 (the agriculture and food security authority interface), and the cadastral interface is the underlying interface that makes the agriculture interface's products actionable. The authority's reliance on the cadastre join is the reliance that allows the procurement advisories to reach the correct proprietors, and the production forecasts to reflect the correct areas, and the reliance is the reliance that the synchronization audit underwrites.

The energy sector regulator and the infrastructure operator are the third indirect consumers. The regulator's and the operator's infrastructure risk decisions depend on the energy infrastructure risk assessment (Chapter 21), and the assessment's join to the cadastre is the join that allows the regulator and the operator to identify the parcels with the high-risk assets, the parcels' proprietors (for the maintenance access agreements), and the parcels' encumbrances (for the easement enforcement). The regulator's and the operator's interface with the geospatial office is treated in Chapter 26 (the energy sector regulator interface), and the cadastral interface is the underlying interface that makes the energy interface's products actionable.

The commercial minerals customers and the strategic stockpile authority are the fourth indirect consumers. The customers' exploration decisions and the authority's stockpile decisions depend on the mineral prospectivity study (Chapter 22), and the study's join to the cadastre is the join that allows the customers to identify the parcels with the high prospectivity, the parcels' proprietors (for the access negotiations), and the parcels' land use codes (for the land-use constraints on the exploration). The customers' and the authority's interface with the geospatial office is treated in Chapter 27 (the commercial customers interface) and Chapter 28 (the strategic reserve authority interface), and the cadastral interface is the underlying interface that makes the minerals interface's products actionable.

The cadastral authority itself is a fifth consumer, in the sense that the authority consumes the office's discrepancy reports. The reports contribute to the cadastre's maintenance, and the authority's consumption of the reports is the authority's contribution to the interface's two-way character. The authority's reliance on the reports is not total — the authority can maintain the cadastre without the office's reports, although the cadastre will be less current — but the authority's acknowledgement of the reports is the acknowledgement that maintains the interface's reciprocity. Prompt and accurate discrepancy reporting is the contribution the office makes to the cadastre's maintenance: the reports allow the cadastral authority to correct the errors the office discovers, and the corrections improve the cadastre's currency for all the consumers, including the office itself. The interface is a two-way relationship: the cadastral authority provides the cadastre, and the geospatial office provides the discrepancy reporting that helps the cadastre stay current.

23.6 The Cadence

The cadastral authority interface operates on four cadences: the daily feed ingestion, the weekly synchronization audit, the monthly discrepancy reporting cycle, and the annual framework review. The four cadences are the rhythms of the interface, and each is maintained by a different role in the office.

The daily feed ingestion is the data pipeline operator's responsibility, conducted as the first task of the morning shift. The operator initiates the ingestion pipeline (the WFS GetFeature request, the REST API call, or the file download), monitors the pipeline's execution (the data is fetched, validated, transformed, and loaded without errors), and verifies the ingestion's completion (the office's cadastre layer reflects the cadastral authority's current record). The daily ingestion is the cadence that ensures the office's cadastre layer is never more than 24 hours behind the cadastral authority's record, and the cadence is the cadence that allows the weekly synchronization audit to identify the changes by comparing the office's layer against the change log for the previous week. The daily ingestion is logged in the ingestion log, which records the date, the time, the source, the record count, and the operator's verification.

The weekly synchronization audit is the data pipeline operator's and the technical lead's joint responsibility, conducted on the first morning of the working week. The operator conducts the audit's operational steps (the change detection, the change application, the topology check, the attribution check), and the technical lead reviews the results and authorizes the updates. The audit's procedures are specified in Chapter 8, section 8.4, and the job breakdown sheet in Chapter 15, section 15.3. The weekly audit is the cadence that ensures the office's cadastre layer is consistent with the cadastral authority's record, and the cadence is the cadence that allows the office to produce the cadastre-dependent products with the confidence that the join is against the current cadastre. The audit is documented in the synchronization audit log, and the log is the audit trail that demonstrates the office's due diligence to the consumers and to the oversight bodies.

The monthly discrepancy reporting cycle is the office manager's responsibility, conducted in the first week of each month. The cycle comprises the discrepancy register's compilation (the discrepancies identified during the previous month's analytical work are compiled into the register), the discrepancy reports' submission (the register's entries are submitted to the cadastral authority through the formal discrepancy reporting channel), and the resolution tracking (the previously

submitted discrepancies' resolutions are tracked, and the resolved discrepancies' updates are applied to the office's cadastre layer). The cycle is the cadence that ensures the office's contribution to the cadastre's maintenance is timely and systematic, and the cadence is the cadence that maintains the interface's two-way character. The cycle is documented in the discrepancy reporting register, and the register is the record that demonstrates the office's contribution to the cadastre's quality.

The annual framework review is the office manager's responsibility, conducted in the first quarter of each calendar year. The review comprises the framework's terms' review (the cadastre feed's format and update frequency, the discrepancy reporting's channel and turnaround time, the synchronization audit's procedures, and the escalation path are reviewed against the current standards and technologies), the framework's update (the terms that need updating are negotiated with the cadastral authority and incorporated into the framework), and the framework's re-signing (the updated framework is signed by the office manager and the cadastral authority's counterpart, and the signed framework is filed with the national spatial data infrastructure's coordination body). The annual review is the cadence that ensures the framework remains current and effective, and the cadence is the cadence that maintains the interface as an institutional arrangement. The review is documented in the framework's review record, and the record is the audit trail that demonstrates the office's institutional stewardship of the interface.

The four cadences' coordination is the office manager's responsibility. The daily ingestion is conducted by the operator, the weekly audit is conducted by the operator and the technical lead, the monthly cycle is conducted by the office manager, and the annual review is conducted by the office manager. The four cadences run in parallel, and each cadence's results inform the next: the daily ingestion's logs are the input to the weekly audit, the weekly audit's discrepancy findings are the input to the monthly cycle, the monthly cycle's resolution tracking is the input to the next month's audit, and the annual review's framework updates shape the next year's ingestion, audit, and cycle. The four cadences' coordination is the operational discipline that ensures the interface is maintained as an ongoing institutional arrangement, not as a one-time setup that erodes over time.

23.7 The Classification Regime

The cadastral authority interface operates under a tiered classification regime: the cadastre layer itself is Internal, the cadastre-dependent analytical products inherit the classification of the analytical content, and the discrepancy reports are Internal with the cadastral authority as the authorized recipient. The tiered regime reflects the cadastre's mixed sensitivity: the parcel geometry is public information in most jurisdictions (the cadastral authority publishes the parcel map for public inspection), but the ownership information is personal data protected by the privacy legislation, and the encumbrance information may be commercially sensitive.

The cadastre layer itself is Internal. The layer contains the parcel geometry, the parcel identifiers, the area, the address, the land use code, the encumbrances, the easements, and the ownership information (the title number, the registered proprietor, the tenure type). The geometry and the basic attribution (the identifier, the area, the address, the land use code) are public information in most jurisdictions, and the cadastral authority publishes them through the public cadastre portal. The ownership information, however, is personal data protected by the privacy legislation (the General Data Protection Regulation in the European Union, the Privacy Act in Australia, the Personal

Data Protection Act in Singapore, the various state public records laws in the United States), and the disclosure of the ownership information is restricted to the authorised purposes (the property transactions, the taxation, the legal proceedings, and the government analytical work). The Internal classification reflects the mixed sensitivity: the layer is available to the office's personnel for the analytical work, but it is not published with the analytical products, and it is not disclosed to the consumers without the cadastral authority's authorisation.

The cadastre-dependent analytical products inherit the classification of the analytical content. The flood inundation map (Chapter 19) carries the tiered classification (Public summary, Internal operational map, Security Confidential impact assessment) determined by the flood domain's classification regime, and the cadastre join is reflected in the product's metadata (the cadastre's version date and the synchronization audit's reference) but not in the product's content (the parcel identifiers are shown on the Internal operational map and the Security Confidential impact assessment, but not on the Public summary). The crop condition assessment (Chapter 20) carries the split classification (Public crop summaries, Internal procurement advisories) determined by the food security domain's classification regime, and the cadastre join is reflected in the Internal procurement advisories (the parcel identifiers and the registered proprietors are shown for the procurement planning) but not in the Public crop summaries (the assessments are aggregated to the district level, with no parcel-level information). The energy infrastructure risk assessment (Chapter 21) carries the Security Confidential classification determined by the energy domain's classification regime, and the cadastre join is reflected in the assessment (the parcel identifiers, the registered proprietors, and the encumbrances are shown for the maintenance access and the easement enforcement). The mineral prospectivity study (Chapter 22) carries the mixed classification (Public base geology, Confidential commercial commissions, Security Confidential critical minerals inventory) determined by the minerals domain's classification regime, and the cadastre join is reflected in the commercial commissions and the critical minerals inventory (the parcel identifiers and the registered proprietors are shown for the access negotiations and the stockpile planning) but not in the public base geology (the geology is published without the parcel join).

The discrepancy reports are Internal, with the cadastral authority as the authorised recipient. The reports contain the parcel identifier, the discrepancy's description, the evidence, and the recommended action, and the reports are submitted to the cadastral authority through the formal discrepancy reporting channel. The Internal classification reflects the reports' content: the parcel identifier and the registered proprietor are personal data, the discrepancy's description may include the alleged error's detail (which could be commercially sensitive if the discrepancy affects a transaction in progress), and the recommended action may include the suggested resurvey (which could prejudice the boundary dispute's resolution if disclosed to the disputing parties). The cadastral authority is the authorised recipient because the authority is the statutory body responsible for the cadastre's maintenance, and the authority's investigation and resolution of the discrepancies is the statutory process that the reports inform.

The classification regime's enforcement is the production operator's responsibility, conducted through the classification marking verification procedure (Chapter 16, section 16.3) before the products' release. The verification checks every product for the correct marking: the cadastre layer's Internal marking (the layer is marked Internal in the metadata and in the access control system), the

cadastre-dependent products' inherited marking (the products carry the analytical content's classification, with the cadastre join reflected in the metadata), and the discrepancy reports' Internal marking with the cadastral authority as the authorised recipient (the reports are marked Internal and addressed to the cadastral authority's discrepancy reporting channel). The office manager's review (Chapter 13, section 13.4) verifies the marking before the products' release, and the review's verification is the final check that ensures the classification regime's integrity. A product that is misclassified has compromised either the public's right to the public information, or the proprietors' privacy, or the customers' commercial confidentiality, or the jurisdiction's strategic security, and the compromise is the failure that the classification regime is designed to prevent.

23.8 The Comparative Practice

The comparative practice of cadastre integration falls into two principal models: the integrated land administration model and the separated-authority model. The two models produce different interface geometries, and the geospatial office's relationship with the cadastral authority is shaped by the model the jurisdiction has adopted. The two models are not the only models — the hybrid model (the cadastre and the land registry in one authority, the national mapping in another) is also practised — but the two principal models are the models that define the comparative practice.

The Singapore Land Authority (SLA), established in 2001 through the merger of the Land Office, the Singapore Land Registry, the Survey Department, and the Lands and Properties Department, is the exemplar of the integrated land administration model. The SLA maintains the cadastre (the parcel boundaries, the survey control network, the Singapore Land Title strata system), the land registry (the title registrations, the mortgages, the easements), and the national mapping function (the base map, the topographic map, the satellite imagery archive) in a single institution. The SLA's Integrated Land Information Service (INLIS) provides the integrated land information to the government's analytical agencies and to the public through a single interface, and the Singapore Geospatial Collaborative Environment, established under the national spatial data infrastructure's coordination, provides the inter-agency data-sharing framework. The integrated model simplifies the geospatial office's interface: the office's interface is with a single authority that provides the cadastre, the mapping, and the land registry, and the interface is conducted under a single inter-agency framework. The model's advantage is the reduced inter-agency coordination cost and the improved data consistency; the model's disadvantage is the concentration of the institutional power in a single authority, which may create its own conflicts of interest (the authority that maintains the cadastre also determines the boundaries, registers the transactions, and produces the base map, and the separation of these functions in different authorities is the safeguard that prevents the concentration).

The Netherlands' Cadastre, Land Registry and Mapping Agency (Kadaster), established in 1832 as the Cadastre of the Netherlands and reorganized in 1995 as the Kadaster, is another exemplar of the integrated land administration model. The Kadaster maintains the cadastre (the parcel boundaries, the survey control network), the land registry (the title registrations, the mortgages, the easements), and the national mapping function (the topographic map, the large-scale base map, the elevation model) in a single institution. The Kadaster's distribution platform (the PDOK — Publieke Dienstverlening Op de Kaart — the public geodata platform) provides the integrated land information to the government's analytical agencies and to the public through a single interface, and the Dutch Spatial

Planning Act and the Land Information Council provide the inter-agency data-sharing framework. The Kadaster's role as the national INSPIRE contact point ensures the cadastre's compliance with the European Union's INSPIRE Directive's Annex II Cadastral Parcels theme. The integrated model's advantages and disadvantages are the same as in Singapore: the reduced coordination cost and the improved data consistency, balanced against the concentration of the institutional power.

The United States' separated-authority model is the exemplar of the separated-authority model in a federation. The cadastre is a county-level function (the county recorder or the county assessor maintains the parcel map and the ownership records for the private lands), the national mapping is a federal function conducted by the United States Geological Survey (the topographic map, the hydrography, the elevation, the land cover), and the public lands cadastre is a federal function conducted by the Bureau of Land Management (the Public Land Survey System, the federal land parcels, the mineral estate). The geospatial office's interface is with multiple authorities — the county recorders for the private lands, the Bureau of Land Management for the federal lands, and the USGS for the national mapping — and the interface requires coordination across the levels of government. The model's advantage is the distributed institutional power and the local accountability (the county recorder is accountable to the county's voters, and the federal mapping agency is accountable to the federal oversight bodies); the model's disadvantage is the increased coordination cost and the data inconsistency (the parcel maps' formats, the attribution's schemas, and the projection's conventions vary across the 3,000-plus county recorders, and the national aggregation of the cadastre — attempted by the Federal Geographic Data Committee's Cadastral Subcommittee through the PLSS CadNSDI dataset — is a substantial and ongoing effort).

The French model is the exemplar of the hybrid model. The Direction Générale des Finances Publiques (DGFiP), the tax authority, maintains the cadastre administratif (the parcel map and the ownership records for the property taxation), and the Service de la Publicité Foncière, housed in the DGFiP since 2017, maintains the land registry (the title registrations, the mortgages, the easements). The Institut Géographique National (IGN, now the Institut national de l'information géographique et forestière), established in 1940 and reorganized in 2012 as a public administrative establishment under the Ministry of Ecology, conducts the national mapping function (the topographic map, the large-scale base map, the orthoimagery). The geospatial office's interface is with two authorities — the DGFiP for the cadastre and the land registry, the IGN for the national mapping — and the interface requires coordination between the two. The hybrid model's advantage is the partial integration (the cadastre and the land registry are in a single authority, simplifying the property transaction process) and the partial separation (the national mapping is in a separate authority, preserving the analytical independence); the model's disadvantage is the historical mismatch between the cadastre and the land registry (the cadastre administratif is a fiscal document, not a legal document, and the parcel boundaries shown on the cadastre are not the legally binding boundaries, which are determined by the deeds and the Surveyor's plans held by the land registry), which the 2017 reorganization has not fully resolved.

These four cases — Singapore, the Netherlands, the United States, and France — confirm the interface's doctrine. The cadastre feed is the operational mechanism in all four cases: the SLA's INLIS, the Kadaster's PDOK, the county recorders' parcel maps distributed through the state geographic information officers' portals, and the DGFiP's cadastre distributed through the IGN's

Géoportail. The synchronization audit is the operational check in all four cases: the office's cadastre layer is compared against the cadastral authority's published change log, and the discrepancies are resolved. The discrepancy reporting is the constructive contribution in all four cases: the office reports the discrepancies discovered through the analytical work, and the cadastral authority resolves them through the statutory process. The inter-agency framework is the institutional arrangement in all four cases: the Singapore Geospatial Collaborative Environment, the Dutch Spatial Planning Act and the Land Information Council, the Federal Geographic Data Committee's Cadastral Subcommittee, and the French Comité national de l'information géographique (CNIG). The model determines the institutional complexity of the interface (the integrated model has one authority, the hybrid has two, the separated has multiple), but the functional boundary is the same in all four cases, and the boundary is the boundary between recording and analysing.

23.9 Worked Example: A Synchronization Discrepancy at the Riverbourne Basin

In the third week of October, the data pipeline operator conducted the weekly cadastre synchronization audit and identified a discrepancy in the Riverbourne Basin, the same district that had experienced the spring flood event treated in Chapters 1, 7, 10, 11, 13, 14, 15, 17, and 19. The operator's audit compared the office's cadastre layer (last synchronized on the previous Monday, 14 October) against the cadastral authority's published change log for the week ending 20 October. The change log identified 47 modified parcels in the Riverbourne Basin, of which 42 were routine subdivisions and consolidations and 5 were boundary adjustments. The operator applied the 42 routine changes without incident, but the 5 boundary adjustments raised a concern: the changes were concentrated in the reach of the Riverbourne River between the Northbridge and the Southbridge, the same reach where the spring flood inundation map (Chapter 19) had identified the floodplain boundary.

The operator's concern was that the cadastre's boundary adjustments might reflect the river's migration since the previous cadastre update, and the migration might have implications for the flood inundation map's accuracy. The operator documented the concern in the synchronization audit log and escalated it to the technical lead. The technical lead reviewed the audit's results, retrieved the flood inundation map produced during the spring flood event, and compared the map's inundation extent against the updated cadastre's parcel boundaries. The comparison revealed that 12 parcels that had been shown as partially inundated on the spring flood map were now split by the updated boundary: the portion of the parcel on the river side of the new boundary had been transferred to the adjacent parcel, and the portion on the land side had been retained. The spring flood map, joined to the now-outdated cadastre, showed the inundation on the wrong parcel.

The technical lead assessed the implications. The spring flood map had been delivered to the emergency management authority on 15 April, and the authority had used the map for the evacuation notifications, the damage assessments, and the relief programmes. The cadastre's boundary adjustments were dated 18 October, six months after the spring flood event, and the adjustments reflected the river's migration that had occurred during the spring flood itself (the flood's erosion and deposition had shifted the river's channel, and the riparian proprietors had subsequently agreed to the boundary adjustments to reflect the new channel). The implication was that the spring flood map had

been correct at the time of its production (the cadastre at that time reflected the pre-flood boundary, and the inundation was correctly attributed to the then-registered parcels), but the map was now incorrect if used for the current work (the parcels' boundaries had changed, and the inundation attribution would be wrong if the map were re-joined to the current cadastre).

The technical lead documented the assessment in the synchronization audit log and notified the office manager. The office manager reviewed the assessment and determined the appropriate response: the spring flood map's archival version (the version delivered to the emergency management authority on 15 April, joined to the cadastre as it stood on that date) was retained unchanged, with a note in the catalog record indicating that the cadastre had subsequently been adjusted and that the map's parcel attribution reflected the pre-adjustment cadastre. The note is the provenance annotation that ensures the archival map's integrity: a consumer who retrieves the archival map in the future will see that the parcel attribution reflects the pre-adjustment cadastre, and will not be misled into thinking that the map's attribution reflects the current cadastre. The office manager also determined that the spring flood map, if re-produced for the current analytical work (such as the post-flood damage assessment or the floodplain re-zoning), would be re-joined to the current cadastre, and the re-joined map would be a new product with a new version date and a new provenance record.

The operator, the technical lead, and the office manager also identified a discrepancy in the cadastre itself. The updated cadastre's boundary adjustments reflected the river's migration, but the river's migration had been documented in the office's flood model's post-event analysis (Chapter 17, section 17.5): the model had identified the river's migration through the comparison of the pre-flood and the post-flood satellite imagery, and the migration's extent had been quantified (the channel had shifted by an average of 8 metres, with a maximum of 14 metres at the Northbridge reach). The cadastral authority's boundary adjustments reflected the migration, but the adjustments' documentation did not reference the office's post-event analysis, and the adjustments' surveyed positions differed from the model's quantified positions by an average of 2 metres (the surveyed positions were based on the field survey conducted by the riparian proprietors' surveyors, and the field survey's accuracy was subject to the survey's methodology and the monumentation's precision). The office manager determined that the discrepancy should be reported to the cadastral authority through the formal discrepancy reporting channel.

The office manager compiled the discrepancy report. The report comprised the parcel identifiers for the 5 adjusted parcels, the discrepancy's description (the boundary adjustments' surveyed positions differed from the office's flood model's quantified positions by an average of 2 metres), the evidence (the office's post-event flood model's analysis, with the satellite imagery comparison, the channel shift quantification, and the model's accuracy assessment), the recommended action (a review of the boundary adjustments' surveyed positions against the flood model's quantified positions, with a possible resurvey if the differences exceeded the survey's tolerance), and the office's contact for the follow-up. The report was submitted to the cadastral authority through the formal discrepancy reporting channel on 25 October, and the submission was logged in the discrepancy reporting register.

The cadastral authority acknowledged the report on 28 October and initiated the investigation. The authority's surveyors reviewed the boundary adjustments' surveyed positions against the office's

flood model's quantified positions, and the review confirmed the differences (the average 2-metre difference, with the maximum 3-metre difference at the Northbridge reach). The authority determined that the differences were within the survey's tolerance (the cadastral survey's standard tolerance for the rural parcels is plus or minus 0.30 metres, but the riparian boundaries are subject to the accretion and the avulsion doctrine, which allows for the natural shift of the watercourse without the boundary adjustment's resurvey, and the 2-metre difference was within the doctrine's tolerance). The authority notified the office of the determination on 12 November, and the office's cadastre layer was confirmed as consistent with the cadastral authority's record. The discrepancy report's resolution was documented in the discrepancy reporting register, and the resolution's notification was filed with the synchronization audit log.

This worked example illustrates the cadastral authority interface's full method: the synchronization audit (the weekly comparison of the office's cadastre layer against the cadastral authority's change log, the identification of the 47 modified parcels in the Riverbourne Basin, the application of the 42 routine changes and the escalation of the 5 boundary adjustments), the analytical implication (the technical lead's comparison of the spring flood map against the updated cadastre, the identification of the 12 parcels whose attribution was affected by the boundary adjustments, the assessment of the map's archival integrity and the current-work implications), the institutional response (the office manager's determination of the archival map's retention with the provenance annotation and the current-work map's re-joining to the current cadastre), and the discrepancy reporting (the office manager's compilation of the report, the submission through the formal channel, the cadastral authority's investigation and determination, and the resolution's documentation). The interface's integrity is the integrity that the synchronization audit ensures, the analytical product's integrity is the integrity that the provenance annotation preserves, and the cadastre's maintenance is the maintenance that the discrepancy reporting contributes to. The interface is the foundation of the office's cadastre-dependent intelligence delivery, and the interface's integrity is the condition of the office's usefulness to the consumers who rely on the cadastre join.

Chapter 24 — The Emergency Management Authority Interface

24.1 The Interface Defined

The emergency management authority interface is the boundary between the office that produces flood risk intelligence and the authority that conducts the operational response. The geospatial office produces the flood hazard maps, the flood inundation maps, the exposure reports, the evacuation guidance, and the event-specific situational awareness products. The emergency management authority receives the intelligence, makes the operational decisions — the evacuation orders, the resource deployment, the public warnings, the mutual-aid requests — and conducts the response. The two functions are operationally interdependent because the authority's decisions require the office's intelligence, and they are institutionally distinct because the analytical independence to produce intelligence and the operational authority to conduct a response are different integrity requirements that should not be combined in a single office.

The interface is governed by the intelligence-response division. When the office assumes operational command of the response, it exceeds its mandate and compromises its analytical independence, because an office conducting the response is tempted to produce intelligence that supports its own decisions rather than intelligence that informs the authority's decisions. When the authority assumes the production of intelligence, it exceeds its mandate and compromises its operational focus, because an authority producing intelligence is tempted to interpret the intelligence in favour of its preferred operational course rather than to weigh the intelligence on its merits. The division is the structural condition that allows both offices to perform their functions with the focus their functions require, and it is maintained through the activation procedure (Chapter 13, section 13.6), the standing inter-agency framework, and the post-event review (section 24.8 below).

The interface is bilateral in its primary geometry: the geospatial office and the emergency management authority are the two parties. The interface is multilateral in its operational expression, because the authority coordinates the response with the police, the fire service, the health authority, the public works department, the transportation authority, and the military or civil-defence assets that may be deployed. The office's intelligence is delivered to the authority, and the authority distributes the intelligence to the operational responders as the authority's operational plan requires. The office does not deliver intelligence directly to the operational responders, because the direct delivery would bypass the authority's coordination and would risk producing inconsistent operational pictures among the responders. The authority is the office's single point of contact for the response, and the office's intelligence is delivered through the authority's common operating picture.

The interface's institutional location varies across jurisdictions. In some jurisdictions, the geospatial office sits within the same ministry as the emergency management authority (a "vertical" interface); in others, the two offices sit in different ministries or different levels of government (a "horizontal" interface). The vertical interface simplifies the administrative coordination but introduces a risk that the analytical independence is compromised by the ministry's operational priorities. The horizontal interface preserves the analytical independence but requires more formal coordination mechanisms. The interface's geometry does not change the functional division: the

office produces the intelligence, and the authority conducts the response, regardless of whether the two offices sit in the same ministry or in different ministries.

24.2 The Scope

The interface covers five activities: standing intelligence provision, event-driven intelligence production, situational awareness sharing, post-event review, and framework maintenance.

Standing intelligence provision is the production and delivery of the office's standing flood risk products — the flood hazard maps, the flood inundation maps for the design return periods, the exposure reports, and the evacuation guidance — to the authority on the standing cadence. The standing products are produced by the office's hazard specialist through the hydrology plug-in (Chapter 10, section 10.2) and the production operator (Chapter 16), and they are delivered to the authority through the secure data interface. The authority uses the standing products for the preparedness phase: the evacuation planning (the identification of the evacuation routes, the assembly points, the shelters), the resource positioning (the pre-positioning of the swift-water rescue teams, the high-axle vehicles, the mobile pumping units), and the public education (the publication of the flood hazard maps for the public awareness). The standing products carry the Public or Internal classification depending on the content: the general flood hazard maps are Public (intended for the public awareness), and the operational flood inundation maps that include the detailed exposure data are Internal (restricted to the authority and the operational responders).

Event-driven intelligence production is the production and delivery of the event-specific intelligence during a flood event. The event-specific intelligence comprises the event-specific inundation map (the map of the current or forecast inundation extent, updated as the event evolves), the event-specific exposure report (the population, the buildings, the critical infrastructure, and the hazardous-materials sites within the inundation extent), the event-specific evacuation guidance (the recommended evacuation zones, the evacuation routes, and the shelter capacities), and the event-specific situational awareness products (the satellite imagery, the social-media reports, and the field reports that the office synthesizes into a common operating picture). The event-driven intelligence is produced on the event's timeline, which may require updates every 6 to 12 hours during a major event, and is delivered to the authority through the secure data interface and the direct communication with the office manager.

Situational awareness sharing is the provision of the office's current analytical picture to the authority's emergency operations centre during an event. The sharing is conducted through the office's production dashboard, which the authority's emergency operations centre may access to view the office's current intelligence picture, including the latest inundation map, the latest exposure report, the latest satellite imagery, and the latest model runs. The situational awareness sharing is a read-only view: the authority may view the office's intelligence but may not modify it, and the office retains the analytical control of the intelligence. The sharing is governed by the access control, which restricts the dashboard access to the authority's authorized personnel and logs every access for the audit trail.

Post-event review is the joint review of the event response after the event is resolved. The review is conducted by the office and the authority jointly, with each office presenting its perspective on the event, the response, and the intelligence support. The review covers the timeline (when the office

produced each intelligence product, when the authority received it, when the authority acted on it), the products delivered (which products were used, which were not used, which were used but found inadequate), the issues encountered (the data gaps, the model limitations, the communication failures, the coordination gaps), the lessons learned, and the corrective actions. The review is documented, and the documentation is the basis for the improvement of the interface, the products, and the procedures for the next event. The post-event review is treated in detail in section 24.8 below.

Framework maintenance is the standing maintenance of the inter-agency framework between events. The framework comprises the data-sharing agreement (which specifies the data the office provides, the formats, the delivery channels, the access controls), the communication protocol (which specifies the communication channels, the contact persons, the escalation procedures), the exercise and simulation programme (which specifies the joint exercises and simulations conducted between events to test the interface), and the review cycle (which specifies the periodic review of the framework's effectiveness). The framework is maintained by the office manager, who is the office's point of contact for the authority, and the framework is reviewed annually or after every major event, whichever comes first.

The scope does not include the operational command of the response, which is the authority's function. It does not include the public communication, which is the authority's function (the office does not issue public warnings or public advisories; the office produces the intelligence, and the authority communicates the intelligence to the public). It does not include the resource deployment, which is the authority's function (the office does not direct the swift-water rescue teams or the high-axle vehicles; the office produces the intelligence that informs the deployment, and the authority deploys the resources). The office that expands its scope to include the command, the communication, or the deployment exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

24.3 The Method

The method of the interface is the structured exchange of intelligence products between the office and the authority, governed by the inter-agency framework and executed through the standing and event-driven procedures. The method comprises the product formats, the delivery channels, the communication procedures, the escalation procedures, and the joint exercises.

The product formats are determined by the authority's operational systems. The authority's emergency operations centre typically operates a common operating picture platform — a web-based GIS that aggregates the intelligence from the multiple sources and displays it on a common map. The office's products must be delivered in the formats that the platform ingests: GeoTIFF for the raster products (the inundation maps, the satellite imagery), Shapefile or GeoPackage for the vector products (the evacuation zones, the critical infrastructure points), KML for the products that may be viewed in the virtual globe viewers, and PDF for the products that may be printed and read by the decision-makers. The product formats are specified in the data-sharing agreement, and the office's production operator produces the products in the agreed formats. The format specifications include the coordinate reference system (the authority's common operating picture platform typically operates in the national or regional projection, such as the Web Mercator for the web-based platforms or the national projection for the desktop GIS), the metadata (the ISO 19115 metadata profile,

including the product's creation date, the producer, the classification, the data sources, and the processing steps), and the symbology (the standard symbology for the flood inundation maps: the blue gradient for the inundation depth, the red for the evacuation zones, the yellow for the shelters, the green for the evacuation routes).

The delivery channels are determined by the inter-agency framework. The standing products are delivered through the secure data interface, which is a bidirectional web service that allows the office to push the products to the authority's platform and the authority to pull the products from the office's archive. The interface operates over the HTTPS protocol with the mutual TLS authentication, ensuring that the products are transmitted securely and that only the authorized systems may access the interface. The interface is monitored for the availability and the performance, and the office's data pipeline operator is responsible for the interface's technical operation. The event-driven products are delivered through the secure data interface for the routine products and through the direct communication with the office manager for the urgent products. The direct communication may be conducted by the phone (the office manager's mobile phone, with the number provided to the authority's emergency operations centre), by the secure messaging (the encrypted messaging application approved for the inter-agency communication), or by the physical presence (the office manager or a designated liaison officer deployed to the authority's emergency operations centre during a major event). The direct communication is reserved for the urgent products that cannot wait for the routine delivery, and the direct communication is logged in the office's event log for the audit trail.

The communication procedures govern the routine and the event-driven communication between the office and the authority. The routine communication is conducted through the periodic coordination meetings (typically monthly, with the office manager and the authority's preparedness chief), the joint exercises (typically twice a year, with the office's hazard specialist and the authority's emergency operations centre staff), and the informal communication (the office manager's periodic check-ins with the authority's preparedness chief). The event-driven communication is conducted through the activation procedure (the office manager activates the office's response upon the authority's request, treated in Chapter 13, section 13.6), the situational awareness sharing (the office's production dashboard, accessible to the authority's emergency operations centre), and the product delivery (the secure data interface and the direct communication). The communication procedures specify the contact persons (the office manager is the office's primary point of contact; the authority's preparedness chief is the authority's primary point of contact for the routine communication, and the authority's emergency operations centre director is the primary point of contact for the event-driven communication), the contact channels (the phone numbers, the email addresses, the secure messaging accounts), and the response times (the office manager responds to the authority's activation request within 30 minutes during the business hours and within 2 hours outside the business hours).

The escalation procedures govern the escalation of the issues that cannot be resolved at the operational level. The issues may include the data gaps (the office's intelligence does not cover the area of the authority's concern), the model limitations (the office's model does not represent the event's specific characteristics), the communication failures (the secure data interface is down, or the direct communication is not reaching the authority), and the coordination gaps (the office's intelligence is not being used by the authority's operational responders, or the authority's operational

responders are producing their own intelligence in parallel with the office's). The escalation procedures specify the escalation path: the operational issue is escalated from the office manager and the authority's preparedness chief to the office's parent ministry and the authority's parent ministry, and the escalation is documented in the office's event log. The escalation is reserved for the issues that the operational level cannot resolve, and the escalation is the mechanism that ensures the interface's problems are addressed at the level that has the authority to address them.

The joint exercises and simulations are the method's standing test. The exercises are conducted twice a year, with the office and the authority jointly simulating a flood event and exercising the interface's procedures. The exercises test the secure data interface's operation, the direct communication's reachability, the product formats' compatibility, the response times, and the coordination between the office's production and the authority's operational response. The exercises are followed by the after-action reviews, which identify the issues encountered and the corrective actions. The exercises are the method's standing test, and the test is the mechanism that ensures the interface is ready when the next event occurs. An office that does not exercise the interface discovers the interface's failures during the event, when the failures have consequences; an office that exercises the interface discovers the failures during the exercise, when the failures can be corrected without consequences.

24.4 The Products

The interface's products are the office's flood risk intelligence products, delivered to the authority on the standing and event-driven cadence. The products comprise five types, each serving a different phase of the emergency management cycle.

The flood hazard map is the standing product, produced on the annual or biennial cadence. The map shows the flood hazard zones for the design return periods (the 10-year, 50-year, 100-year, and 500-year flood hazard zones, derived from the hydrological and hydraulic modelling treated in Chapter 19, section 19.3). The map is used by the authority for the preparedness phase: the evacuation planning (the identification of the evacuation routes and the assembly points), the resource positioning (the pre-positioning of the rescue teams and the equipment), the public education (the publication of the hazard zones for the public awareness), and the land-use planning consultation (the authority's input to the planning authority's decisions on the development in the flood hazard zones). The map carries the Public classification for the general version (intended for the public awareness) and the Internal classification for the operational version (which includes the additional detail on the depths, the velocities, and the arrival times that are restricted to the authority and the operational responders).

The flood inundation map is the event-driven product, produced during a flood event. The map shows the current or forecast inundation extent, updated as the event evolves. The current inundation map is derived from the satellite imagery (the Sentinel-1 radar imagery, which can penetrate the cloud cover that typically accompanies the flood events, or the Sentinel-2 optical imagery, when the cloud cover permits), the hydraulic model's simulation of the current conditions (the model run with the observed rainfall and the observed river stages), and the field reports (the authority's field teams' observations of the inundation extent). The forecast inundation map is derived from the hydraulic model's forecast run (the model run with the forecast rainfall and the forecast river stages, producing

the forecast inundation extent for the next 12 to 48 hours). The map is used by the authority for the response phase: the evacuation orders (the identification of the areas to be evacuated), the resource deployment (the dispatch of the rescue teams to the inundated areas), and the public warnings (the issuance of the warnings to the affected populations). The map carries the Internal classification by default, and the Security Confidential classification if the inundation extent affects the critical infrastructure (the substations, the water-treatment plants, the hazardous-materials sites) whose geolocation is security-sensitive.

The exposure report is the analytical product, produced both on the standing cadence (the annual or biennial exposure report based on the flood hazard map) and on the event-driven cadence (the event-specific exposure report based on the flood inundation map). The report identifies the population, the buildings, the critical infrastructure, and the hazardous-materials sites within the inundation extent, and quantifies the exposure: the number of people exposed, the number of buildings exposed (categorized by the building type: residential, commercial, industrial, public), the critical infrastructure exposed (categorized by the type: substations, water-treatment plants, hospitals, schools, fire stations, police stations, transportation hubs), and the hazardous-materials sites exposed (the chemical storage facilities, the fuel depots, the industrial sites with the hazardous materials). The report is used by the authority for the response phase: the prioritization of the evacuation (the areas with the highest population exposure are evacuated first), the resource allocation (the rescue teams are dispatched to the areas with the highest building exposure), and the critical infrastructure protection (the utility operators are notified of the infrastructure exposure, and the hazardous-materials teams are dispatched to the hazardous-materials sites). The report carries the Internal classification by default, and the Security Confidential classification if it includes the critical infrastructure's geolocation.

The evacuation guidance is the operational product, produced during a flood event. The guidance identifies the recommended evacuation zones (derived from the inundation map, with the zones defined by the clear geographical boundaries: the roads, the rivers, the administrative boundaries), the evacuation routes (the roads that remain passable during the flood event, derived from the road network's flood vulnerability assessment), and the shelter capacities (the designated shelters' locations and their current capacities, derived from the authority's shelter registry). The guidance is used by the authority for the evacuation orders: the authority's emergency operations centre uses the guidance to issue the evacuation orders, to direct the evacuation traffic, and to manage the shelter operations. The guidance carries the Internal classification, and the Public classification for the version that is communicated to the public (which excludes the operational detail on the route capacities and the shelter vulnerabilities).

The post-event assessment is the analytical product, produced after the flood event is resolved. The assessment comprises the event's timeline (the rainfall, the river stages, the inundation extent, and the response actions, in chronological order), the event's impact (the population affected, the buildings damaged, the infrastructure damaged, the economic loss), the response's effectiveness (the evacuation's timeliness, the rescue's coverage, the shelter's adequacy), and the intelligence support's effectiveness (which intelligence products were used, which were not used, which were used but found inadequate, and the recommendations for the improvement). The assessment is used by the

authority for the recovery phase and the preparedness phase for the next event. The assessment carries the Internal classification, with the Public summary version for the public communication.

24.5 The Consumers

The interface's primary consumer is the emergency management authority itself. The authority's emergency operations centre is the immediate consumer during an event: the centre's director and the operations officers use the flood inundation map, the exposure report, and the evacuation guidance to make the operational decisions — the evacuation orders, the resource deployment, the public warnings, the mutual-aid requests. The authority's preparedness division is the consumer between events: the division's planners use the flood hazard map, the standing exposure report, and the post-event assessments to develop the evacuation plans, the resource positioning plans, and the public education programmes. The authority's recovery division is the consumer after the event: the division's case managers use the post-event assessment to manage the recovery assistance programmes (the temporary housing, the infrastructure repair, the business recovery loans).

The authority's operational responders are the indirect consumers, receiving the office's intelligence through the authority's common operating picture. The police receive the evacuation guidance for the traffic control and the area security. The fire service receives the inundation map and the exposure report for the swift-water rescue and the property protection. The health authority receives the exposure report for the hospital evacuation and the medical response planning. The public works department receives the inundation map for the road closures and the infrastructure protection. The transportation authority receives the evacuation guidance for the public transport mobilization. The military or civil-defence assets, when deployed, receive the full intelligence picture for the support operations. The office does not deliver the intelligence directly to the operational responders; the authority distributes the intelligence through its common operating picture, and the distribution is the authority's function.

The public is the ultimate beneficiary of the interface's intelligence, but the public is not the office's direct consumer. The public receives the intelligence through the authority's public communication: the flood hazard maps (published by the authority for the public awareness), the evacuation orders (issued by the authority during the event), the shelter information (published by the authority during the event), and the post-event recovery information (published by the authority after the event). The office's intelligence underpins the authority's public communication, but the office does not communicate directly with the public, because the public communication during an emergency is the authority's function, and the authority's communication must be consistent with the authority's operational decisions. The office that communicates directly with the public during an emergency risks producing inconsistent messages that undermine the public's trust in the authority's response.

The supporting agencies are the secondary consumers. The meteorological service provides the rainfall forecasts that the office uses as the input to the hydraulic model, and the meteorological service may receive the office's flood forecasts in return for the verification of the meteorological forecasts' accuracy. The hydrological service (where it is separate from the office) provides the river stage observations and forecasts that the office uses as the input and the validation of the hydraulic model. The cadastre authority provides the parcel data that the office uses for the exposure analysis

(Chapter 23). The statistical office provides the demographic data that the office uses for the population exposure analysis. The national mapping agency (where it is separate from the office) provides the base map data that the office uses for the product cartography. The supporting agencies' contributions to the office's intelligence are governed by the bilateral data-sharing arrangements, and the office's intelligence is delivered to the supporting agencies only when the authority's common operating picture includes them as the authorized recipients.

The political leadership is the tertiary consumer, receiving the office's intelligence through the authority's briefings. The political leadership (the elected officials, the senior appointed officials) receives the event briefings from the authority's emergency operations centre during an event, and the office's intelligence is summarized in the briefings. The political leadership uses the briefings to make the policy decisions (the declaration of the emergency, the mobilization of the additional resources, the request for the federal or international assistance). The office does not brief the political leadership directly during an event, because the political leadership's briefings must be consistent with the authority's operational picture, and the authority's emergency operations centre is the single point of contact for the political leadership's briefings. The office may brief the political leadership directly between events, when the briefings are the standing intelligence products (the flood hazard maps, the post-event assessments) rather than the event-driven intelligence, and the briefings are conducted through the office's parent ministry.

24.6 The Cadence

The interface operates on three cadences: the standing cadence between events, the event-driven cadence during events, and the post-event cadence after events.

The standing cadence comprises the monthly coordination meeting (the office manager and the authority's preparedness chief review the office's recent products, the authority's recent operations, the upcoming events, and the framework's operation), the biannual joint exercise (the office and the authority jointly simulate a flood event, exercising the interface's procedures), the annual flood hazard map update (the office produces the updated flood hazard map, incorporating the new data and the new modelling), the annual exposure report update (the office produces the updated exposure report, incorporating the new demographic and the new infrastructure data), and the annual framework review (the office and the authority jointly review the framework's effectiveness, identifying the issues encountered and the corrective actions). The standing cadence is maintained by the office manager, who is the office's point of contact for the authority, and the cadence is the mechanism that maintains the interface's readiness between events.

The event-driven cadence is activated when a flood event occurs. The activation is initiated by the authority's request (the authority's emergency operations centre contacts the office manager to request the event-driven intelligence) or by the office's initiative (the office's monitoring detects the event's onset, and the office manager contacts the authority to offer the event-driven intelligence). The activation triggers the office's event-driven response (Chapter 13, section 13.6), which comprises the immediate production of the event-specific intelligence (the event-specific inundation map, exposure report, and evacuation guidance, produced within 2 to 4 hours of the activation), the situational awareness sharing (the office's production dashboard, accessible to the authority's emergency operations centre), and the ongoing production (the regular updates of the event-specific

intelligence, every 6 to 12 hours during the event's evolution). The event-driven cadence is maintained by the hazard specialist, who produces the event-specific intelligence, the production operator, who produces and delivers the products, and the office manager, who coordinates with the authority's emergency operations centre and authorizes the releases.

The event-driven cadence's intensity scales with the event's severity. A minor flood event (the river stage exceeding the minor flood stage but not the moderate flood stage) may require only the standing products' confirmation and the situational awareness sharing. A moderate flood event (the river stage exceeding the moderate flood stage but not the major flood stage) requires the event-specific intelligence production on the 12-hour update cycle. A major flood event (the river stage exceeding the major flood stage) requires the event-specific intelligence production on the 6-hour update cycle, the deployment of the office manager or the designated liaison officer to the authority's emergency operations centre, and the activation of the office's surge capacity (the recall of the off-duty staff, the reassignment of the staff from the other domains). The intensity's scaling is specified in the activation procedure, and the office manager determines the intensity in consultation with the authority's emergency operations centre director.

The post-event cadence is activated when the event is resolved. The post-event cadence comprises the immediate post-event review (conducted within 2 weeks of the event's resolution, with the office and the authority jointly reviewing the event's timeline, the products delivered, the issues encountered, and the immediate lessons learned), the detailed post-event assessment (produced within 3 months of the event's resolution, with the office producing the comprehensive assessment of the event's impact and the intelligence support's effectiveness), and the corrective action implementation (the implementation of the corrective actions identified in the review and the assessment, with the implementation tracked through the monthly coordination meetings until completion). The post-event cadence is maintained by the office manager and the hazard specialist, who produce the post-event assessment, and the office manager, who tracks the corrective actions' implementation.

The cadence is the mechanism that maintains the interface's readiness and the interface's improvement over time. The standing cadence maintains the readiness between events. The event-driven cadence delivers the intelligence during events. The post-event cadence improves the interface after events. The three cadences together constitute the interface's operational rhythm. The rhythm is what allows a small office to maintain the interface with a large authority, despite the asymmetry in the two offices' sizes and the infrequency of the major events that test the interface.

24.7 The Classification Regime

The interface operates under a tiered classification regime, with the classification determined by the product's content and the product's intended recipients. The tiered regime comprises three levels: Public, Internal, and Security Confidential.

The Public classification applies to the general flood hazard maps and the public summaries of the post-event assessments. These products are intended for the public awareness, and they do not include the operational detail that could compromise the response or the security-sensitive information that could inform the attacks on the critical infrastructure. The Public flood hazard map shows the flood hazard zones (the 100-year and the 500-year flood hazard zones, without the detailed

depths, velocities, and arrival times), the general evacuation guidance (the designated evacuation routes and the shelter locations, without the operational detail on the capacities and the vulnerabilities), and the general exposure information (the number of people and buildings in the flood hazard zones, without the parcel-level detail). The Public classification allows the authority to publish the products on its website, to distribute them through the media, and to use them in the public education programmes. The office's production operator produces the Public versions of the products, and the office manager authorizes the Public release after verifying that the products do not include the operational or the security-sensitive information.

The Internal classification applies to the operational flood inundation maps, the detailed exposure reports, the operational evacuation guidance, and the detailed post-event assessments. These products are intended for the authority and the operational responders, and they include the operational detail that the public versions exclude. The Internal flood inundation map shows the inundation extent with the depths and the arrival times, the detailed exposure report shows the parcel-level exposure, and the operational evacuation guidance shows the route capacities and the shelter vulnerabilities. The Internal classification restricts the products to the authority's authorized personnel and the operational responders, and the products are delivered through the secure data interface with the access control. The office's production operator applies the Internal marking, and the office manager reviews the marking before the release.

The Security Confidential classification applies to the event-specific intelligence that includes the critical infrastructure's geolocation or vulnerabilities. During a flood event, the inundation extent may affect the critical infrastructure (the substations, the water-treatment plants, the hospitals, the hazardous-materials sites), and the event-specific intelligence that identifies the affected infrastructure carries the Security Confidential classification. The classification is consistent with the critical infrastructure information protection frameworks (treated in Chapter 21, section 21.7), and it restricts the products to the security-cleared personnel in the authority, the operational responders, and the intelligence agencies. The Security Confidential products are delivered through the secure data interface with the additional access controls (the role-based access control, the need-to-know basis, the access logging), and the products are stored in the office's classified archive (Chapter 12, section 12.5).

The classification regime's tiered structure allows the office to produce the intelligence at the appropriate level of restriction for each audience. The Public products inform the public awareness. The Internal products inform the operational response. The Security Confidential products inform the critical infrastructure protection. The classification regime is maintained by the production operator, who applies the markings, and the office manager, who reviews the markings before the release. The regime is non-negotiable: every product carries the appropriate classification, and the marking is applied in the standard format (capital letters, top and bottom of every page, in the metadata) and in the standard location. The office manager may declassify a product only if the product's content does not include the operational detail or the critical infrastructure information, and the declassification is documented in the production log with the rationale. The classification regime is the mechanism that allows the office to deliver the intelligence to the multiple audiences — the public, the operational responders, and the security-cleared personnel — without compromising the response or the critical infrastructure's protection.

24.8 The Comparative Practice

The comparative practice of the emergency management authority interface falls into three principal models: the unified command model, the coordinated agency model, and the embedded liaison model. The three models produce different interface geometries, and the geospatial office operates under whichever model the jurisdiction has adopted.

The unified command model places the geospatial office and the emergency management authority under a unified command structure during an event. The model is exemplified by the United States Federal Emergency Management Agency's National Incident Management System (NIMS), established in 2004 under Homeland Security Presidential Directive 5. NIMS establishes the Incident Command System (ICS) as the standard command structure for the major incidents, with the unified command allowing the multiple jurisdictions and the multiple agencies to manage the incident jointly. Under the unified command, the geospatial office's intelligence is delivered through the ICS Planning Section's Situation Unit, which is responsible for the common operating picture. The geospatial office may deploy a liaison officer to the Emergency Operations Centre (EOC) or the Incident Command Post (ICP) as a Technical Specialist under the Situation Unit. The unified command model simplifies the coordination during the event, because the intelligence is delivered through the established command structure, but the model may compromise the geospatial office's analytical independence if the office is placed under the operational command of the authority. The model is appropriate for the jurisdictions where the geospatial office is part of the emergency management authority's parent department (such as the geospatial functions within the United States Department of Homeland Security).

The coordinated agency model maintains the geospatial office and the emergency management authority as separate agencies that coordinate during an event. The model is exemplified by the European Union's Copernicus Emergency Management Service (CEMS), operated by the European Commission's Joint Research Centre. The CEMS provides the satellite-based intelligence to the national civil protection authorities of the EU member states and to the European Commission's Emergency Response Coordination Centre (ERCC), without assuming the operational command. The CEMS's Rapid Mapping component provides the satellite-based damage assessment within hours of an activation, using the Sentinel imagery and the commercial high-resolution imagery. The activation is requested by the authorised users (the national civil protection authorities and the ERCC) through the CEMS portal, and the products are delivered through the same portal. The CEMS does not deploy liaison officers to the national civil protection authorities' emergency operations centres; the coordination is conducted through the portal and the phone. The coordinated agency model preserves the geospatial office's institutional independence, but the model requires the more formal coordination mechanisms (the activation procedure, the product formats, the delivery channels, the access controls) to compensate for the absence of the unified command. The model is appropriate for the jurisdictions where the geospatial office is institutionally separate from the emergency management authority (such as the EU-level CEMS, which is institutionally separate from the national civil protection authorities).

The embedded liaison model places a designated liaison officer from the geospatial office in the emergency management authority's emergency operations centre during an event. The model is exemplified by the Geoscience Australia's relationship with the Emergency Management Australia

(EMA, within the Department of Home Affairs). Geoscience Australia, as the national geoscience agency, produces the earthquake, flood, and other hazard intelligence, and during a major event, Geoscience Australia deploys a liaison officer to the EMA's Crisis Coordination Centre (CCC). The liaison officer provides the direct interpretation of the geospatial intelligence, advises on the additional products that may be required, and facilitates the rapid tasking of the satellite imagery. The embedded liaison model combines the geospatial office's institutional independence (the liaison officer remains a Geoscience Australia employee) with the close coordination of the unified command (the liaison officer is physically present in the authority's emergency operations centre). The model is appropriate for the jurisdictions where the geospatial office and the emergency management authority are separate agencies but where the events' severity warrants the close coordination.

The Japanese model is a fourth variation, with the Cabinet Office's Disaster Management Bureau coordinating the response and the Geospatial Information Authority of Japan (GSI, within the Ministry of Land, Infrastructure, Transport and Tourism) providing the geospatial intelligence. GSI operates the emergency mapping system that produces the flood, earthquake, and volcanic hazard maps, and during a major event, GSI activates the Disaster Response Headquarters and provides the intelligence to the Cabinet Office's Emergency Disaster Response Headquarters. The Japanese model is a coordinated agency model with the strong Cabinet-level coordination, and the model reflects the Japanese disaster management system's emphasis on the central coordination of the major events. The 2011 Tohoku earthquake and tsunami demonstrated the model's strength: GSI produced the rapid damage maps within hours of the event, the Cabinet Office coordinated the response, and the geospatial intelligence was used to direct the rescue and the relief operations across the affected prefectures.

The four models share the intelligence-response division: the geospatial office produces the intelligence, and the emergency management authority conducts the response. The models differ in the institutional location of the coordination (the unified command under NIMS, the portal-based coordination under CEMS, the liaison officer under Geoscience Australia, the Cabinet-level coordination under GSI), but the functional division is the same in all models. The geospatial office operates under whichever model the jurisdiction has adopted, and the office's interface with the emergency management authority is shaped by the model. The model determines the institutional complexity of the interface, but the interface's substantive content — the products, the delivery channels, the communication procedures, the classification regime — is the same across the models, and the substantive content is the subject of this chapter.

24.9 Worked Example: The Riverbourne Basin Flood Event Interface

The Riverbourne Basin, treated in Chapters 1, 7, 10, 11, 13, 14, 15, 17, and 19, experiences a major flood event in late March 2024. The event is triggered by an intense rainfall event on 22-24 March, with the 72-hour rainfall totalling 215 millimetres across the basin (exceeding the 100-year, 72-hour rainfall depth of 185 millimetres derived in Chapter 19, section 19.9). The Riverbourne at Portsmith reaches a peak stage of 4.2 metres above the flood stage on 25 March, exceeding the previous record of 3.52 metres set in 2019 (a 50-year event). The event inundates approximately 18

square kilometres of the Portsmith metropolitan area, affecting an estimated 38,000 residents and damaging an estimated 4,200 buildings.

The interface's activation begins at 09:15 on 23 March, when the meteorological service issues a severe weather warning for the Riverbourne Basin. The office's data pipeline operator, who monitors the meteorological service's feeds as part of the morning ingestion check (Chapter 15, section 15.6), notes the warning and notifies the office manager. The office manager contacts the emergency management authority's preparedness chief at 09:30 to inform the authority of the warning and to offer the office's event-driven intelligence. The preparedness chief acknowledges the warning and requests the office's event-specific intelligence production to begin. The office manager activates the office's event-driven response at 09:45, recalling the hazard specialist from leave and reassigning the production operator from the standing production to the event-driven production.

The event-specific intelligence production begins at 10:00 on 23 March. The hazard specialist configures the hydrology plug-in for the event-specific run: the HEC-HMS model is configured with the observed rainfall (from the rain gauge network and the radar rainfall estimates) and the forecast rainfall (from the meteorological service's numerical weather prediction model, which forecasts an additional 80 millimetres over the next 24 hours). The HEC-RAS 2D model is configured with the LiDAR DEM (the same 2-metre, 0.12-metre RMSE DEM used in the 100-year flood hazard map, treated in Chapter 19, section 19.9), the updated Manning's n values, and the boundary conditions (the forecast hydrograph from HEC-HMS for the upstream boundary, the forecast storm surge level from the meteorological service for the downstream boundary). The model run takes approximately 90 minutes and produces the forecast inundation extent for the next 48 hours at 2-metre grid resolution.

The hazard specialist produces the first event-specific inundation map at 11:45 on 23 March. The map shows the forecast inundation extent for 24 March at 06:00 (the forecast peak), with the depths colour-coded by the blue gradient (0-0.5 metres, 0.5-1.0 metres, 1.0-2.0 metres, 2.0 metres and above). The map identifies the inundation of approximately 12 square kilometres of the Portsmith metropolitan area, including the central business district's low-lying areas, the residential neighbourhoods along the Riverbourne's left bank, and the industrial estate at the basin's outlet. The specialist produces the event-specific exposure report at 12:30, identifying approximately 28,000 residents and 3,100 buildings within the forecast inundation extent, including 2 hospitals, 4 schools, 1 substation, 1 water-treatment plant, and 2 hazardous-materials sites. The specialist produces the event-specific evacuation guidance at 13:00, identifying 6 recommended evacuation zones, 8 evacuation routes, and 5 shelters with the combined capacity of 4,500 people.

The production operator produces the products in the agreed formats: GeoTIFF for the inundation map (for ingestion into the authority's common operating picture platform), Shapefile for the evacuation zones (for the authority's GIS), KML for the products that may be viewed in the virtual globe viewers by the authority's field teams, and PDF for the products that may be printed and read by the authority's decision-makers. The operator applies the classification markings: the inundation map is Internal (with the Security Confidential version produced separately, identifying the affected critical infrastructure), the exposure report is Internal (with the Security Confidential version produced separately, identifying the hospitals, the substation, the water-treatment plant, and the hazardous-materials sites by name and address), and the evacuation guidance is Internal (with the Public version produced separately, excluding the operational detail on the route capacities and the

shelter vulnerabilities). The operator delivers the products through the secure data interface at 13:30, and the office manager contacts the authority's emergency operations centre director by phone to confirm the delivery and to offer the direct interpretation.

The authority's emergency operations centre, activated at 10:00 on 23 March, receives the products and integrates them into the common operating picture. The centre's director reviews the products at 14:00 and issues the evacuation orders for the 6 recommended evacuation zones at 14:30. The authority's operational responders (the police for the traffic control, the fire service for the swift-water rescue, the public works department for the road closures) are deployed based on the products. The office's intelligence is distributed to the operational responders through the common operating picture, and the office does not communicate directly with the operational responders. The office manager deploys to the authority's emergency operations centre at 16:00 as the liaison officer, providing the direct interpretation of the intelligence and advising on the additional products that may be required.

The event-driven cadence continues through 23-26 March. The hazard specialist produces the updated event-specific intelligence every 6 hours: the updated inundation map (incorporating the observed rainfall and the observed river stages), the updated exposure report (incorporating the updated inundation extent), and the updated evacuation guidance (incorporating the updated evacuation zones and the shelter capacities). The production operator produces and delivers the products on the 6-hour cycle. The office manager, deployed at the authority's emergency operations centre, provides the direct interpretation and facilitates the additional product requests. The office's surge capacity is activated on 24 March, with the recall of the off-duty staff and the reassignment of the staff from the other domains, to maintain the 6-hour cycle through the event's duration.

The event peaks on 25 March at 04:00, when the Riverbourne at Portsmith reaches the peak stage of 4.2 metres above the flood stage. The observed inundation extent is approximately 18 square kilometres, exceeding the forecast extent of 12 square kilometres by 50 per cent (the forecast underestimated the inundation because the observed rainfall of 215 millimetres exceeded the forecast rainfall of 195 millimetres, and because the storm surge coincided with the riverine flood peak, producing the combined inundation that the model's separate boundary conditions did not fully capture). The office's hazard specialist notes the underestimation and produces the revised inundation map at 06:00 on 25 March, incorporating the observed data and the revised forecast. The revised map is delivered to the authority's emergency operations centre at 06:30, and the authority adjusts the evacuation orders and the resource deployment accordingly.

The event recedes on 26-28 March, and the authority stands down the emergency operations centre on 28 March. The office stands down the event-driven response on 28 March, and the office manager returns to the office. The post-event cadence begins: the immediate post-event review is conducted on 5 April, with the office and the authority jointly reviewing the event's timeline, the products delivered, the issues encountered (the underestimation of the inundation extent, the 6-hour cycle's strain on the office's surge capacity, the communication gap between the office's liaison officer and the operational responders), and the immediate lessons learned. The detailed post-event assessment is produced by the office on 15 June, with the comprehensive analysis of the event's impact (38,000 residents affected, 4,200 buildings damaged, $480 million in economic loss) and the intelligence support's effectiveness (the products were used by the authority for the evacuation

orders, the resource deployment, and the public warnings; the underestimation of the inundation extent was the primary issue, and the corrective action is the improvement of the hydraulic model's coupled boundary conditions to better represent the coincidence of the riverine flood peak and the storm surge peak). The corrective actions are tracked through the monthly coordination meetings, and the improved hydraulic model is operational by the 2025 flood season.

This worked example illustrates the interface's full operation: the activation (the meteorological warning at 09:15, the office's notification of the authority at 09:30, the activation at 09:45), the event-driven intelligence production (the HEC-HMS and HEC-RAS runs, the inundation map at 11:45, the exposure report at 12:30, the evacuation guidance at 13:00, the product delivery at 13:30), the authority's operational response (the evacuation orders at 14:30, the resource deployment, the common operating picture), the office manager's liaison deployment (at 16:00 on 23 March), the 6-hour update cycle through 23-26 March, the underestimation and the correction on 25 March, the stand-down on 28 March, and the post-event review and the assessment on 5 April and 15 June. The interface's intelligence-response division was maintained throughout: the office produced the intelligence, the authority conducted the response, and the two functions were coordinated through the secure data interface, the direct communication, and the liaison officer. The interface's classification regime was maintained: the Internal products for the authority and the operational responders, the Security Confidential products for the critical infrastructure information, and the Public products for the public communication. The interface's cadence was maintained: the event-driven cadence during the event, the post-event cadence after the event, and the corrective actions' implementation through the standing cadence. The interface is the foundation of the office's flood risk intelligence delivery to the emergency management authority, and the interface's integrity is the condition of the office's contribution to the public safety during the flood events.

Chapter 25 — The Strategic Reserve Authority Interface

25.1 The Interface Defined

The strategic reserve authority interface is the boundary between the office that produces the procurement advisory and the authority that makes the release and buy decisions for the strategic reserve. The geospatial office produces the weekly crop condition bulletin, the seasonal yield forecast, and the procurement advisory — the three products treated in Chapter 20 — that synthesize the satellite-derived crop condition with the reserve's stock levels to produce a release or buy signal. The strategic reserve authority receives the advisory, decides whether to release grain from the reserve, buy grain for the reserve, or hold, and conducts the procurement. The two functions are operationally interdependent because the authority's decisions require the office's forward-looking intelligence, and they are institutionally distinct because the analytical independence to forecast the yield and the procurement authority to commit public funds are different integrity requirements that should not be combined in a single office.

The interface is governed by the forward-sensor relationship and the intelligence-decision division. The forward-sensor relationship is what distinguishes the geospatial office from a crop-monitoring service. A crop-monitoring service produces crop condition assessments; a forward-sensor for a strategic reserve authority produces procurement advisories that synthesize the crop condition with the reserve's stock levels to produce a decision signal. The intelligence-decision division is the structural condition that allows the office to produce the forward-looking intelligence without assuming the decision-making function. When the office participates in the release or buy decision, it is tempted to produce intelligence that supports the preferred decision; when the authority participates in the intelligence production, it is tempted to request intelligence that supports the preferred decision. The division prevents both temptations, and the division is maintained through the data-sharing arrangement, the advisory's classification protection, and the decision feedback loop treated in sections 25.4 through 25.7 below.

The interface's institutional location varies across jurisdictions. In some jurisdictions, the geospatial office sits within the same ministry as the strategic reserve authority (typically the agriculture ministry or the food ministry); in others, the office sits in a separate ministry (a science ministry, a planning ministry, or a national mapping agency) and the authority sits in the agriculture or food ministry. The interface may also span levels of government, with the office at the national level and the reserve authority at the sub-national level, or vice versa. The institutional geometry does not change the functional division: the office produces the advisory, and the authority decides, regardless of whether the two offices sit in the same ministry or in different ministries.

The interface is bilateral in its primary geometry: the geospatial office and the strategic reserve authority are the two parties. The interface is multilateral in its operational expression, because the authority's decisions affect and are affected by the agricultural ministry (which sets the agricultural policy and may provide the planted area data), the trade ministry (which sets the import and export rules and may provide the trade flow data), the finance ministry (which approves the procurement budget), and the meteorological service (which provides the weather data). The office's advisory is delivered to the authority, and the authority coordinates the procurement decision with the other ministries. The office does not coordinate directly with the other ministries on the procurement

decision, because the coordination is the authority's function, and the office's role is to produce the intelligence that informs the authority's coordination.

25.2 The Scope

The interface covers six activities: standing intelligence provision, procurement advisory production, stock level monitoring, decision feedback reception, framework maintenance, and post-season review.

Standing intelligence provision is the production and delivery of the office's standing food security products — the weekly crop condition bulletin and the seasonal yield forecast — to the authority on the standing cadence. The standing products are produced by the office's resource economics specialist through the crop model plug-in (Chapter 10, section 10.3) and the production operator (Chapter 16), and they are delivered to the authority through the internal distribution channel. The authority uses the standing products for the reserve's ongoing management: the weekly monitoring of the crop conditions that will determine the next season's supply, the seasonal planning of the reserve's stock rotations, and the medium-term planning of the reserve's strategic posture. The standing products carry the Public classification for the crop condition summaries (intended for the public awareness and the agricultural sector) and the Internal classification for the operational versions (which include the detailed yield forecasts and the stock-to-consumption ratios that are restricted to the authority and the authorized ministries).

Procurement advisory production is the production and delivery of the procurement advisory, the domain's forward-sensor product. The advisory is produced as the yield forecast and the stock levels warrant, typically at the seasonal milestones (the planting forecast, the heading forecast, and the harvest forecast, treated in Chapter 20, section 20.4). The advisory synthesizes the yield forecast with the stock levels and the supply-demand analysis to produce the release or buy signal. The advisory carries the Internal classification by default, because its premature disclosure could move market prices. The advisory production is treated in detail in section 25.4 below.

Stock level monitoring is the office's ongoing retrieval and verification of the reserve's stock levels through the data-sharing arrangement with the authority. The stock levels are the critical input to the supply-demand analysis: the supply-demand balance is computed as Production + Reserve Stocks – Consumption, and the reserve stocks' level determines whether the balance is positive (a surplus, recommending a release) or negative (a deficit, recommending a buy). The stock level retrieval is conducted weekly, with the authority providing the current stock levels through the secure data interface. The office verifies the stock levels against the authority's published reports (where available) and against the historical patterns (a sudden change in the stock levels may indicate a release or a buy that the authority has not yet communicated, and the office queries the authority for the explanation).

Decision feedback reception is the office's reception of the authority's decisions and the decisions' outcomes. The feedback comprises the decision (release, buy, or hold), the decision's rationale (which may include considerations beyond the office's advisory), the decision's implementation (the procurement contract, the release schedule, the buy schedule), and the decision's outcome (the effect on the reserve's stock levels and the market's prices). The feedback is the

mechanism that allows the office to assess the advisory's usefulness, and the feedback is treated in detail in section 25.6 below.

Framework maintenance is the standing maintenance of the data-sharing arrangement and the inter-agency framework between the seasonal cycles. The framework comprises the data-sharing agreement (which specifies the stock level data the authority provides, the formats, the delivery channels, the access controls), the advisory production protocol (which specifies the advisory's format, the production cadence, the classification marking, the delivery channel), the decision feedback protocol (which specifies the feedback's format, the timing, the content), and the review cycle (which specifies the post-season review of the advisory's accuracy and the authority's decisions). The framework is maintained by the office manager, who is the office's point of contact for the authority, and the framework is reviewed annually, after each harvest, and revised as the lessons from the season warrant.

Post-season review is the joint review of the season's intelligence and decisions after the harvest is completed and the procurement cycle is closed. The review is conducted by the office and the authority jointly, with each office presenting its perspective on the season's crop conditions, the yield forecast's accuracy, the advisory's usefulness, the authority's decisions, and the decisions' outcomes. The review covers the forecast errors (the difference between the forecast yield and the reported yield, by district and by region), the advisory's signals (whether the signals were correct in direction and magnitude), the authority's decisions (whether the decisions were timely and appropriate), and the lessons learned. The review is documented, and the documentation is the basis for the improvement of the advisory, the forecast, and the framework for the next season. The post-season review is treated in detail in section 25.8 below.

The scope does not include the procurement decision, which is the authority's function. It does not include the reserve management (the physical storage, the rotation, the quality control), which is the authority's function. It does not include the agricultural policy (the production targets, the input subsidies, the trade policy), which is the agricultural ministry's function. It does not include the market analysis (the futures prices, the trade flows, the commodity speculation), which is the market authority's function. The office that expands its scope to include the decision, the management, the policy, or the market analysis exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

25.3 The Method

The method of the interface is the structured production and exchange of the procurement advisory between the office and the authority, governed by the data-sharing arrangement and the advisory production protocol. The method comprises the stock level retrieval, the supply-demand analysis, the advisory compilation, the classification marking, the delivery, and the decision feedback.

The stock level retrieval is conducted weekly through the secure data interface. The authority provides the current stock levels in the agreed format (typically a CSV or an XML file with the stock levels by commodity, by storage facility, and by quality grade), and the office's data pipeline operator retrieves the file and ingests it into the office's database. The retrieval is monitored for the availability and the timeliness: if the file is not received by the agreed time (typically Monday 09:00), the operator

notifies the office manager, who contacts the authority's data liaison to resolve the delay. The stock level data is stored in the office's classified archive (Chapter 12, section 12.5), with the access restricted to the resource economics specialist, the office manager, and the technical lead. The stock level data is the sensitive input to the supply-demand analysis, and its confidentiality is maintained through the access control and the audit logging.

The supply-demand analysis is conducted by the resource economics specialist using the crop model plug-in (Chapter 10, section 10.3) and the supply-demand worksheet. The analysis begins with the yield forecast, which is the crop model plug-in's output: the forecast yield by district and by region, with the confidence interval. The production estimate is computed as Production = Yield × Planted Area, where the Planted Area is from the cadastre or from the agricultural ministry's planted area survey. The consumption estimate is computed as Consumption = Population × Per-Capita Consumption Rate, where the Per-Capita Consumption Rate is derived from the historical consumption data, adjusted for trends (population growth, dietary changes, income growth). The supply-demand balance is computed as Balance = Production + Reserve Stocks – Consumption. The balance's sign and magnitude determine the advisory's signal: a positive balance indicates a surplus and recommends a release; a negative balance indicates a deficit and recommends a buy; a balance near zero indicates a parity and recommends a hold. The signal's magnitude is calibrated to the balance's magnitude: a small surplus may recommend a small release, while a large surplus may recommend a large release; a small deficit may recommend a small buy, while a large deficit may recommend a large buy. The calibration is specified in the advisory production protocol, and the calibration's parameters (the threshold for the release, the threshold for the buy, the recommended quantities) are reviewed annually in the post-season review.

The advisory compilation is conducted by the resource economics specialist using the advisory template. The advisory comprises four sections: the yield forecast section (the current forecast, the change from the previous forecast, the forecast's confidence interval, the forecast's assumptions), the stock level section (the current stock levels, the stock-to-consumption ratio, the stock levels' trend), the supply-demand analysis section (the production estimate, the consumption estimate, the supply-demand balance, the balance's interpretation), and the signal section (the advisory's conclusion: release, buy, or hold, with the recommended quantity and the rationale). The advisory is compiled in the standard format (a 4-to-6-page PDF with the standard header, the classification marking, and the advisory's reference number), and the compilation is conducted in the office's advisory production environment, which is separate from the public-facing production environment to maintain the advisory's confidentiality.

The classification marking is applied by the production operator before the advisory's release. The advisory carries the Internal classification by default, because its premature disclosure could move market prices. The release signal, if disclosed before the authority's release decision, would allow market participants to anticipate the release and adjust their positions, which would move prices against the authority's interest. The buy signal, if disclosed before the authority's buy decision, would allow market participants to anticipate the buy and adjust their positions, which would move prices against the authority's interest. The Internal marking is applied in the standard format (capital letters, top and bottom of every page, in the metadata) and in the standard location. The marking is

non-negotiable, and the operator that applies the marking incorrectly is subject to the corrective action specified in the office's quality management procedure.

The delivery is conducted through the internal distribution channel, which restricts the advisory to the authority's authorized recipients. The channel is a secure email distribution list, maintained by the office manager and the authority's data liaison, with the recipients specified by name and role. The delivery is logged in the office's distribution log, with the date, the time, the recipients, and the advisory's reference number. The delivery is acknowledged by the authority's primary recipient (typically the authority's procurement director), and the acknowledgement is recorded in the log. The advisory that is not acknowledged within 24 hours is followed up by the office manager, who contacts the authority's procurement director to confirm the receipt.

The decision feedback is received from the authority through the same internal distribution channel. The feedback comprises the decision (release, buy, or hold), the decision's rationale, the decision's implementation (the procurement contract, the release schedule, the buy schedule), and the decision's outcome (the effect on the reserve's stock levels and the market's prices). The feedback is received by the office manager, who records it in the office's decision log, and the feedback is reviewed by the resource economics specialist in the next advisory production cycle. The decision feedback is treated in detail in section 25.6 below.

25.4 The Advisory Production

The advisory production is the interface's central activity, and it is the activity through which the office exercises its forward-sensor function. The production is conducted on the seasonal cadence, with the advisories produced at the planting forecast, the heading forecast, and the harvest forecast milestones (treated in Chapter 20, section 20.4), and on the event-driven cadence, with the advisories produced when the crop conditions or the stock levels change significantly between the milestones.

The planting forecast advisory is produced at the beginning of the growing season, typically 4 to 6 weeks after the planting is completed. The advisory's yield forecast is based on the planned area, the historical yield, and the seasonal weather outlook. The advisory's stock levels are the current reserve stocks, obtained from the authority. The advisory's supply-demand analysis is the initial assessment of the season's balance, and the advisory's signal is the initial recommendation: typically a hold (because the season's outcome is still uncertain) or a preliminary buy (if the planned area is significantly below the previous year and the stock-to-consumption ratio is low). The planting forecast advisory's uncertainty is high, and the advisory explicitly states the uncertainty (the confidence interval on the yield forecast, the assumptions on the weather and the area) to prevent the authority from over-relying on the advisory's signal.

The heading forecast advisory is produced at the heading stage, approximately mid-season, when the crop's condition is well-established and the yield forecast's accuracy is significantly improved. The advisory's yield forecast is based on the condition assessment, the weather to date, and the weather outlook for the remaining growing period. The advisory's stock levels are the current reserve stocks. The advisory's supply-demand analysis is the updated assessment, and the advisory's signal is the updated recommendation: the hold may be continued, the preliminary buy may be confirmed or revised, or a release may be recommended (if the condition is above-normal and the surplus is anticipated). The heading forecast advisory is the most consequential of the three seasonal

advisories, because it provides the authority with the reliable signal that the authority needs to time the procurement (the buy contracts for the import, the release schedules for the domestic release) before the harvest.

The harvest forecast advisory is produced approximately two weeks before the expected harvest, when the crop is near maturity and the yield forecast's accuracy is at its highest. The advisory's yield forecast is based on the near-harvest condition and the weather to date. The advisory's stock levels are the current reserve stocks. The advisory's supply-demand analysis is the final assessment, and the advisory's signal is the final recommendation: the release or the buy that the authority should execute after the harvest. The harvest forecast advisory is the most accurate of the three seasonal advisories, and it is the advisory that the authority uses to finalize the procurement decisions for the season.

The event-driven advisories are produced between the milestones when the crop conditions or the stock levels change significantly. A significant change in the crop conditions may be a drought that develops between the planting and the heading, a flood that affects the planted area, a pest outbreak that damages the crop, or a heat wave that reduces the yield potential. A significant change in the stock levels may be a release that the authority executes in response to a price spike, a buy that the authority executes to replenish the reserve, or a quality issue that removes stocks from the reserve (a spoilage, a contamination). The event-driven advisory is produced within 48 hours of the significant change's detection, and the advisory's signal is the updated recommendation that accounts for the change. The event-driven advisory's production is treated in the worked example in section 25.9 below.

The advisory production is conducted by the resource economics specialist, who compiles the advisory; the technical lead, who validates the forecast and the supply-demand analysis; and the office manager, who reviews the advisory before the release. The production is conducted in the office's advisory production environment, which is separate from the public-facing production environment, and the advisory's draft is stored in the office's classified archive with the access restricted to the production team. The advisory's release is authorized by the office manager, and the release is logged in the office's distribution log. The advisory's archive copy is retained for 10 years (the standard retention period for the procurement advisories, aligned to the audit cycle of the authority's procurement decisions).

25.5 The Products

The interface's products are the office's food security intelligence products, delivered to the authority on the standing, seasonal, and event-driven cadence. The products comprise five types, each serving a different phase of the reserve management cycle.

The weekly crop condition bulletin is the standing product, produced every week during the growing season and delivered by close of business each Wednesday. The bulletin's crop condition assessment section presents the condition map (the green-to-red condition gradient, with the parcel boundaries overlaid), the condition statistics (the percentage of the planted area in each condition class), and the condition trend. The bulletin's yield forecast update section presents the updated yield forecast (by district and by region, with the forecast's confidence interval) and the forecast's change from the previous update. The bulletin's seasonal progress summary section presents the phenological stage, the weather summary, and the seasonal outlook. The bulletin carries the Public

classification for the general version (intended for the public awareness and the agricultural sector) and the Internal classification for the operational version (which includes the detailed yield forecasts and the stock-to-consumption ratios).

The seasonal yield forecast is the milestone product, produced at the planting, the heading, and the harvest milestones. The forecast presents the yield by district and by region, with the confidence interval, the forecast's assumptions (the weather scenario, the soil moisture, the management conditions), and the forecast's change from the previous milestone. The forecast is the analytical core of the advisory, and the forecast's accuracy is the condition of the advisory's usefulness. The forecast carries the Internal classification, with the Public summary version (which presents the regional yield estimates without the district-level detail) for the public communication.

The procurement advisory is the forward-sensor product, produced at the seasonal milestones and on the event-driven cadence. The advisory comprises the four sections described in section 25.4 above, and the advisory's signal is the recommendation that the authority uses to time the procurement. The advisory carries the Internal classification, and the advisory's distribution is restricted to the authority's authorized recipients. The advisory is the interface's central product, and the advisory's accuracy, timeliness, and confidentiality are the conditions of the authority's trust.

The stock level report is the supporting product, produced weekly by the office based on the stock level data retrieved from the authority. The report presents the current stock levels by commodity, by storage facility, and by quality grade, the stock-to-consumption ratio, the stock levels' trend, and the comparison with the historical patterns. The report is used by the resource economics specialist as the input to the supply-demand analysis, and the report is delivered to the office manager and the technical lead for the awareness. The report carries the Internal classification, and the report is not delivered to the authority (the authority already has its own stock level data).

The post-season assessment is the analytical product, produced after the harvest is completed and the procurement cycle is closed. The assessment comprises the season's crop conditions (the chronology of the condition assessments, the weather's influence, the anomalies), the yield forecast's accuracy (the comparison of the forecast yield with the reported yield, by district and by region, with the error analysis), the advisory's signals (the chronology of the advisories, the signals' direction and magnitude, the signals' correctness), the authority's decisions (the chronology of the decisions, the decisions' implementation, the decisions' outcomes), and the lessons learned. The assessment is used by the authority for the next season's planning, and the assessment is the basis for the post-season review (section 25.8). The assessment carries the Internal classification, with the Public summary version for the public communication.

25.6 The Decision Feedback

The decision feedback is the mechanism by which the authority informs the office of the decisions made on the basis of the office's advisory, and the mechanism by which the office learns whether its advisories are useful. The feedback is the authority's communication to the office, and the office does not evaluate the authority's decisions: the office's function is intelligence, not decision-making, and the authority's decisions are the authority's responsibility.

The feedback comprises four elements: the decision, the rationale, the implementation, and the outcome. The decision is the authority's choice among the three options: release (grain is released

from the reserve to the market), buy (grain is purchased for the reserve from the market or from the imports), or hold (no release and no buy). The decision is communicated to the office within 24 hours of the decision's making, through the internal distribution channel. The rationale is the authority's reasoning, which may include the considerations beyond the office's advisory: the fiscal implications (the cost of the procurement, the budget's availability), the political implications (the public perception of the release or the buy, the electoral considerations), the trade implications (the import contracts, the export restrictions, the trade agreements), and the strategic implications (the reserve's long-term management, the stock-to-consumption ratio's target). The rationale is communicated to the office within 7 days of the decision's making, and the rationale allows the office to understand the considerations that the authority weighed in addition to the office's advisory.

The implementation is the authority's execution of the decision: the procurement contract (the supplier, the quantity, the price, the delivery schedule), the release schedule (the quantity, the timing, the distribution channels), or the buy schedule (the quantity, the timing, the source). The implementation is communicated to the office as it progresses, with the updates on the contract's execution, the delivery's status, and the quality's verification. The implementation is the mechanism that allows the office to track the decision's effect on the reserve's stock levels, and the stock levels' update is the input to the next advisory's supply-demand analysis.

The outcome is the decision's effect on the reserve's stock levels and the market's prices. The stock levels' outcome is the change in the reserve's stocks after the decision's implementation: the release reduces the stocks, the buy increases the stocks, and the hold maintains the stocks. The market prices' outcome is the change in the market prices after the decision's announcement and implementation: the release typically reduces the prices (the increased supply lowers the prices), the buy typically increases the prices (the increased demand raises the prices), and the hold typically maintains the prices. The outcome is communicated to the office within 30 days of the decision's implementation, and the outcome allows the office to assess the advisory's accuracy (whether the forecast yield matched the reported yield), the advisory's signal's correctness (whether the signal's direction matched the decision's direction), and the advisory's timeliness (whether the advisory was produced early enough to inform the decision).

The feedback is reviewed by the resource economics specialist in the next advisory production cycle, and the feedback is the input to the post-season review. The specialist uses the feedback to assess the advisory's strengths and weaknesses: the advisories that were followed by the decisions in the recommended direction are the successes, and the advisories that were followed by the decisions in the opposite direction (or not followed at all) are the failures. The specialist analyzes the failures to identify the causes: the forecast errors (the yield forecast was inaccurate), the signal errors (the signal's direction was wrong), the timeliness errors (the advisory was produced too late), or the considerations beyond the advisory (the authority's fiscal, political, trade, or strategic considerations outweighed the advisory's signal). The causes' identification is the basis for the corrective actions, and the corrective actions are tracked through the post-season review.

The decision feedback is the mechanism that makes the forward-sensor relationship a learning relationship. An office that receives the feedback honestly, acknowledging its advisory's failures as well as its successes, improves its advisory over time. An office that does not receive the feedback, or that receives it defensively, does not improve, and the next season's advisories reproduce the same

failures. The feedback is the mechanism that allows the small office to improve its food security intelligence over time, despite the infrequency of the major supply disruptions that test the advisory's usefulness.

25.7 The Classification Regime

The interface operates under a tiered classification regime, with the classification determined by the product's content and the product's potential to move market prices. The tiered regime comprises two levels: Public and Internal.

The Public classification applies to the crop condition summaries and the public summaries of the seasonal yield forecasts. These products are intended for the public awareness and the agricultural sector, and they do not include the operational detail that could move market prices. The Public crop condition summary presents the regional condition assessments (the percentage of the planted area in each condition class, by region) without the district-level detail. The Public yield forecast summary presents the regional yield estimates (the forecast yield by region, with the confidence interval) without the district-level detail. The Public classification allows the authority to publish the products on its website, to distribute them through the agricultural extension service, and to use them in the public communication. The office's production operator produces the Public versions, and the office manager authorizes the Public release after verifying that the products do not include the operational detail.

The Internal classification applies to the procurement advisory, the operational versions of the weekly crop condition bulletin and the seasonal yield forecast, the stock level report, and the post-season assessment. These products include the operational detail that the public versions exclude: the district-level yield forecasts, the stock-to-consumption ratios, the supply-demand balances, and the release or buy signals. The Internal classification restricts the products to the authority's authorized recipients and the office's authorized personnel, and the products are delivered through the internal distribution channel with the access control. The office's production operator applies the Internal marking, and the office manager reviews the marking before the release.

The Internal classification is the economic integrity guarantee for the procurement advisory. The release signal, if disclosed before the authority's release decision, would allow the market participants to anticipate the release and adjust their positions, which would move the prices against the authority's interest. The buy signal, if disclosed before the authority's buy decision, would allow the market participants to anticipate the buy and adjust their positions, which would move the prices against the authority's interest. The Internal marking is the mechanism that prevents the disclosure, and the marking is non-negotiable. The advisory that is disclosed prematurely has compromised the authority's procurement position, and the compromise may result in the higher procurement costs (for a buy) or the lower release revenues (for a release), with the financial loss borne by the public budget. The office that maintains the Internal marking produces an advisory that is trusted by the authority; an office that does not maintain the marking produces an advisory that is not trusted, and the distrust is the condition that undermines the interface's effectiveness.

The classification regime is consistent with the market-sensitive information protection frameworks in most jurisdictions. The United States Department of Agriculture's World Agricultural Outlook Board, which produces the World Agricultural Supply and Demand Estimates (WASDE)

report, maintains the "lockup" procedure: the report's analysts are sequestered in a secure room from the time the data is compiled until the report's release, with no communication with the outside, to prevent the premature disclosure of the report's estimates. The report is released at the scheduled time (12:00 Eastern Time on the release day), and the lockup ensures that all market participants receive the information simultaneously. The European Commission's Mars bulletin, which provides the crop yield forecasts for the European Union, follows a similar procedure, with the bulletin's release scheduled and the pre-release access restricted. The geospatial office's Internal classification is aligned with these frameworks, and the office's advisory production environment is the equivalent of the lockup: the advisory's draft is produced in the separate environment, the access is restricted, and the release is authorized only at the scheduled time.

The classification regime's tiered structure allows the office to produce the intelligence at the appropriate level of restriction for each audience. The Public products inform the public awareness and the agricultural sector. The Internal products inform the authority's procurement decisions. The classification regime is maintained by the production operator, who applies the markings, and the office manager, who reviews the markings before the release. The regime is non-negotiable: every product carries the appropriate classification, and the marking is applied in the standard format (capital letters, top and bottom of every page, in the metadata) and in the standard location. The office manager may declassify a product (reduce the classification from Internal to Public) only if the product's content does not include the operational detail, and the declassification is documented in the production log with the rationale.

25.8 The Comparative Practice

The comparative practice of the strategic reserve authority interface spans three commodity domains — grain reserves, petroleum reserves, and critical minerals stockpiles — and four institutional models. The grain reserve interface is the principal context for the geospatial office's food security intelligence, and the petroleum and minerals interfaces are the comparative references that illustrate the forward-sensor relationship's applicability across the commodity domains.

The grain reserve interface is exemplified by the Food Corporation of India (FCI) and the National Food Authority of the Philippines (NFA). The FCI, established in 1965 under the Food Corporations Act, maintains India's buffer stock of wheat and rice (approximately 60 million tonnes in 2023, distributed across the FCI's network of storage depots). The FCI's procurement decisions are informed by the Indian Ministry of Agriculture and Farmers Welfare's crop production forecasts, which are produced by the Directorate of Economics and Statistics using the satellite-based crop condition assessments from the Mahalanobis National Crop Forecast Centre (MNCFC). The MNCFC, established in 2012, uses the NDVI and the EVI from the Sentinel-2 and the Resourcesat-2 satellites to produce the crop condition assessments at the district level, and the forecasts are provided to the FCI through the Ministry of Agriculture. The FCI's release and buy decisions are made by the Department of Food and Public Distribution, based on the forecasts and the FCI's stock levels. The interface's forward-sensor relationship is maintained: the MNCFC produces the forecasts, the FCI manages the stocks, and the Department of Food and Public Distribution decides. The NFA, established in 1972 under Republic Act 6 a.k.a. the National Grains Authority Act, maintains the Philippines' buffer stock of rice (approximately 30 days of national consumption). The NFA's

procurement decisions are informed by the Philippine Department of Agriculture's Bureau of Agricultural Research's crop production forecasts, which use the satellite-based assessments from the Philippine Space Agency (PhilSA). The interface's geometry is similar to the FCI's, with the PhilSA producing the forecasts, the NFA managing the stocks, and the Department of Agriculture deciding.

The petroleum reserve interface is exemplified by the United States Strategic Petroleum Reserve (SPR) and the International Energy Agency (IEA) coordinated reserves. The SPR, established in 1975 under the Energy Policy and Conservation Act, maintains the United States' buffer stock of petroleum (approximately 350 million barrels in 2024, stored in four underground salt cavern sites along the Gulf of Mexico). The SPR's release decisions are made by the President of the United States, based on the supply disruption assessments from the Department of Energy's Energy Information Administration (EIA) and the International Energy Agency. The EIA's assessments use the satellite-based monitoring of the global petroleum infrastructure (the tanker traffic, the storage tank fill levels, the refinery operations) from the commercial satellite operators (Maxar, Planet) and from the National Geospatial-Intelligence Agency's EnhancedView programme. The IEA's coordinated reserves require the member states to maintain the reserves equivalent to 90 days of net imports, and the coordinated release decisions are made by the IEA's Governing Board, based on the supply disruption assessments from the IEA's Oil Industry and Markets Division. The interface's forward-sensor relationship is maintained: the geospatial intelligence (from the EIA, the NGA, the commercial operators) produces the supply disruption assessments, the SPR and the IEA member states manage the reserves, and the political leadership decides.

The critical minerals stockpile interface is exemplified by the United States National Defense Stockpile (NDS), managed by the Defense Logistics Agency (DLA) under the Strategic and Critical Materials Stock Piling Act. The NDS maintains the buffer stocks of approximately 50 critical minerals (including the rare earth elements, the lithium, the cobalt, the tungsten, and the antimony), and the release and buy decisions are made by the DLA's Strategic Materials office, based on the supply chain risk assessments from the Department of Defense's Office of Industrial Policy and the Department of the Interior's U.S. Geological Survey (USGS). The USGS's National Minerals Information Center produces the mineral commodity summaries and the supply chain risk assessments, using the satellite-based monitoring of the global mining infrastructure (the active mines, the exploration sites, the processing facilities) from the commercial satellite operators and from the NGA. The interface's forward-sensor relationship is maintained: the USGS produces the supply chain risk assessments, the DLA manages the stockpile, and the political leadership (the President, the Congress) decides on the strategic posture.

The four institutional models observed across the three commodity domains are the integrated model, the separated model, the coordinated model, and the hybrid model. The integrated model places the geospatial office and the reserve authority within the same ministry (the FCI and the MNCFC are both within the Indian government's agricultural and food framework, though in different ministries). The separated model places the two offices in different ministries or agencies (the SPR is within the Department of Energy, and the EIA is also within the Department of Energy, but the NGA is within the Department of Defense). The coordinated model places the two offices in different jurisdictions (the IEA's coordinated reserves involve the national reserve authorities and the national geospatial offices of the member states, coordinated by the IEA's secretariat). The hybrid

model combines the elements of the three models (the NDS involves the DLA, the USGS, and the NGA, with the DLA and the USGS within the U.S. federal framework and the NGA within the intelligence community). The four models share the forward-sensor relationship: the geospatial office produces the advisory, and the reserve authority decides. The models differ in the institutional complexity, but the functional division is the same in all models, and the functional division is the subject of this chapter.

The four models also share the classification challenge. The grain reserve advisory's premature disclosure could move the grain futures prices. The petroleum reserve advisory's premature disclosure could move the petroleum futures prices. The minerals stockpile advisory's premature disclosure could move the minerals spot prices. The Internal classification (or the equivalent in each jurisdiction) is the mechanism that prevents the disclosure, and the classification protection is the shared element of the interface across the commodity domains. The geospatial office operates under whichever model the jurisdiction has adopted, and the office's interface with the reserve authority is shaped by the model. The model determines the commodity, the stockpile mechanism, and the decision criteria, but the forward-sensor relationship, the intelligence-decision division, and the classification protection are the same across the models, and these are the interface's substantive content.

25.9 Worked Example: Producing the Heading Forecast Advisory During the Northfields Drought

The resource economics specialist is producing the heading forecast advisory for the Northfields, the jurisdiction's principal agricultural region treated in Chapter 20, section 20.9. The growing season has been marked by a developing drought: the seasonal rainfall (March to July 2024) is 35 per cent below normal for the Northfields, which produces approximately 60 per cent of the jurisdiction's wheat. The specialist has been tracking the drought's impact on the crop condition through the weekly bulletins, and the condition has deteriorated from "above-normal" in March (when the planting was completed under favourable conditions) to "below-normal" in May (when the rainfall deficit began to affect the crop) to "poor" in July (when the drought intensified and the crop's yield potential was significantly reduced). The bulletin for the week of 15 July 2024, treated in Chapter 20, reported the NDVI at the 22nd percentile of the historical distribution, the EVI at the 25th percentile, and the NDWI at the 18th percentile, with 45 per cent of the planted area in the poor class.

The specialist begins the advisory production on Monday, 22 July 2024, with the stock level retrieval. The data pipeline operator retrieves the current stock level file from the strategic reserve authority through the secure data interface at 09:15. The file shows the reserve's wheat stocks at 1.85 million tonnes, distributed across the authority's four storage depots (the Northfields Depot with 0.65 million tonnes, the Portsmith Depot with 0.50 million tonnes, the Capital Depot with 0.40 million tonnes, and the Western Depot with 0.30 million tonnes). The stock-to-consumption ratio is 3.7 months (the stocks represent 3.7 months of the national wheat consumption, which is approximately 0.50 million tonnes per month), which is below the authority's target of 4.5 months and indicates a low stock level. The specialist notes the low stock-to-consumption ratio and the deteriorating crop condition, and the combination suggests a buy signal for the advisory.

The specialist proceeds with the yield forecast update at 10:00, using the crop model plug-in. The WOFOST crop growth model is configured with the current conditions (the NDVI at the 22nd percentile, the NDWI at the 18th percentile, the soil moisture from the NASA SMAP satellite at 25 per cent below normal), the weather to date (the March-to-July rainfall at 65 per cent of normal, the temperatures at 1.2°C above normal), and the weather outlook for the remaining growing period (the meteorological service's 30-day forecast predicts the below-normal rainfall and the above-normal temperatures to continue). The model run takes approximately 45 minutes and produces the updated yield forecast for the Northfields' wheat: 2.1 tonnes per hectare, with the 90-per-cent confidence interval of 1.8 to 2.4 tonnes per hectare. The forecast is significantly below the previous year's yield of 3.2 tonnes per hectare (a normal year) and below the 5-year average of 2.9 tonnes per hectare. The specialist cross-references the forecast with the statistical regression model (which uses the NDVI at the heading stage and the historical yield), and the regression model produces a forecast of 2.0 tonnes per hectare, which is within the WOFOST model's confidence interval and confirms the forecast's plausibility.

The specialist proceeds with the supply-demand analysis at 11:00, using the supply-demand worksheet. The production estimate is computed as Production = Yield × Planted Area = 2.1 tonnes per hectare × 1.20 million hectares = 2.52 million tonnes. The consumption estimate is computed as Consumption = Population × Per-Capita Consumption Rate = 12 million people × 50 kg per person per year = 0.60 million tonnes per year = 0.050 million tonnes per month. The annual consumption is 0.60 million tonnes, and the carry-over from the previous season's stocks (0.35 million tonnes) is added to the production. The supply is Production + Carry-over = 2.52 + 0.35 = 2.87 million tonnes. The annual supply-demand balance is Balance = Supply – Consumption = 2.87 – 0.60 = 2.27 million tonnes. The balance is positive (a surplus of 2.27 million tonnes), but the surplus must be considered in the context of the reserve's target. The reserve's target is 4.5 months of consumption, which is 4.5 × 0.050 = 0.225 million tonnes, and the surplus is well above the target. However, the current reserve stocks are 1.85 million tonnes (3.7 months of consumption), and the authority's target is to maintain the reserve at 4.5 months. The specialist computes the reserve's projected level after the harvest: the current stocks (1.85 million tonnes) minus the consumption until the harvest (3 months × 0.050 = 0.15 million tonnes) plus the procurement from the harvest (the surplus minus the seed and the feed, approximately 0.40 million tonnes) = 1.85 – 0.15 + 0.40 = 2.10 million tonnes. The projected reserve level is 2.10 million tonnes, which is 4.2 months of consumption — still below the 4.5-month target. The specialist notes that the projected reserve level remains below the target, and the advisory's signal should account for the shortfall.

The specialist compiles the advisory at 12:00, using the advisory template. The yield forecast section presents the updated forecast (2.1 tonnes per hectare, 90-per-cent confidence interval 1.8 to 2.4 tonnes per hectare), the change from the previous forecast (the planting forecast was 2.6 tonnes per hectare, so the heading forecast is a 19-per-cent downward revision), the forecast's assumptions (the WOFOST model with the current conditions and the weather outlook), and the forecast's uncertainty (the confidence interval reflects the drought's intensity and the weather outlook's uncertainty). The stock level section presents the current stocks (1.85 million tonnes, 3.7 months of consumption), the projected stocks after the harvest (2.10 million tonnes, 4.2 months of consumption), and the comparison with the target (4.5 months). The supply-demand analysis section

presents the production estimate (2.52 million tonnes), the consumption estimate (0.60 million tonnes per year), the annual balance (a surplus of 2.27 million tonnes), and the reserve's projected level (2.10 million tonnes, below the 4.5-month target). The signal section presents the advisory's conclusion: a buy of 0.20 million tonnes is recommended to bring the projected reserve level to 2.30 million tonnes (4.6 months of consumption, just above the 4.5-month target), with the rationale that the drought-reduced yield and the below-target reserve level warrant the precautionary buy to maintain the reserve's adequacy through the next consumption cycle. The buy should be executed before the harvest to avoid the price increase that the drought-reduced supply will cause after the harvest.

The specialist submits the advisory draft to the technical lead at 13:00 for the forecast validation. The technical lead reviews the WOFOST model's configuration (the parameters, the inputs, the assumptions), the regression model's calibration (the historical data, the coefficients), and the supply-demand analysis's computations (the production estimate, the consumption estimate, the balance, the reserve's projected level). The technical lead confirms the forecast's plausibility (the WOFOST and the regression models agree within the confidence interval), the analysis's correctness (the computations are accurate), and the signal's appropriateness (the buy recommendation is consistent with the analysis). The technical lead's validation is documented in the advisory's provenance chain.

The specialist submits the validated advisory to the office manager at 14:30 for the release authorization. The office manager reviews the advisory's classification (Internal, applied by the production operator), the advisory's content (the four sections, the signal, the rationale), and the advisory's recipients (the authority's procurement director, the authority's data liaison, the office manager). The office manager authorizes the release at 15:00, and the production operator delivers the advisory through the internal distribution channel at 15:15. The authority's procurement director acknowledges the receipt at 16:30.

The authority's decision is communicated to the office on 25 July 2024, three days after the advisory's delivery. The decision is a buy of 0.20 million tonnes, as recommended by the advisory. The rationale, communicated on 30 July 2024, includes the considerations beyond the advisory: the fiscal implications (the cost of the buy, approximately $80 million at the prevailing international wheat price of $400 per tonne, is within the authority's procurement budget), the trade implications (the import contract is negotiated with a supplier in a neighbouring country, with the delivery scheduled for August-September 2024), and the strategic implications (the buy maintains the reserve's adequacy through the next consumption cycle, and the authority's stock-to-consumption ratio is restored to the 4.5-month target). The implementation is communicated as it progresses: the procurement contract is signed on 5 August 2024, the first delivery arrives on 20 August 2024, and the buy is completed on 15 September 2024. The outcome is communicated on 15 October 2024: the reserve's stocks are increased to 2.05 million tonnes (the buy of 0.20 million tonnes is partially offset by the consumption during the period), and the market prices have increased by 8 per cent since the advisory's release (the drought-reduced supply, confirmed by the harvest report in September 2024, has moved the prices as anticipated).

The decision feedback is reviewed by the specialist in the next advisory production cycle (the harvest forecast, produced in late August 2024). The specialist notes that the advisory's signal was followed by the authority's decision in the recommended direction (the buy of 0.20 million tonnes, as recommended), and the advisory's forecast was accurate (the harvest report, published in September

2024, reports the Northfields' wheat yield at 2.05 tonnes per hectare, which is within the heading forecast's confidence interval of 1.8 to 2.4 tonnes per hectare). The specialist also notes that the advisory's timeliness was appropriate (the advisory was delivered on 22 July 2024, and the authority's decision was made on 25 July 2024, allowing the buy to be executed before the harvest and before the price increase). The advisory is assessed as a success, and the success is documented in the post-season assessment produced in November 2024.

This worked example illustrates the interface's full operation: the stock level retrieval at 09:15 (the secure data interface, the current stocks at 1.85 million tonnes, the stock-to-consumption ratio at 3.7 months), the yield forecast update at 10:00 (the WOFOST model, the forecast yield at 2.1 tonnes per hectare, the confidence interval 1.8 to 2.4 tonnes per hectare, the cross-reference with the regression model), the supply-demand analysis at 11:00 (the production estimate, the consumption estimate, the annual balance, the reserve's projected level below the 4.5-month target), the advisory compilation at 12:00 (the four sections, the buy signal of 0.20 million tonnes, the rationale), the technical lead's validation at 13:00, the office manager's authorization at 15:00, the delivery at 15:15, the authority's decision on 25 July (the buy of 0.20 million tonnes, as recommended), the implementation (the contract signed on 5 August, the delivery completed on 15 September), and the outcome (the reserve's stocks increased, the market prices increased by 8 per cent as anticipated). The interface's forward-sensor relationship was maintained: the office produced the advisory, the authority decided, and the two functions were coordinated through the secure data interface, the internal distribution channel, and the decision feedback. The interface's classification regime was maintained: the Internal marking on the advisory, the restricted distribution, and the access control. The interface's intelligence-decision division was maintained: the office did not participate in the decision, and the authority did not participate in the intelligence production. The interface is the foundation of the office's food security intelligence delivery to the strategic reserve authority, and the interface's integrity is the condition of the office's contribution to the jurisdiction's food security. The drought of 2024 demonstrated the interface's value: the advisory's buy signal, delivered before the harvest and before the price increase, allowed the authority to maintain the reserve's adequacy at a lower cost than would have been possible after the harvest, and the cost saving (approximately $16 million, the 8-per-cent price increase on the 0.20 million tonnes) is the economic value of the office's forward-looking intelligence, and the value is the reason the interface exists.

Chapter 26 — The Sectoral Regulator Interface

26.1 The Interface Defined

The sectoral regulator interface is the boundary between the office that produces intelligence and the regulator that regulates the sector. The geospatial office produces the energy infrastructure intelligence (treated in Chapter 21) and the minerals intelligence (treated in Chapter 22), and the sectoral regulator regulates the energy sector and the minerals sector, holding the asset registry, the licence registry, and the compliance framework. The two functions are operationally interdependent because the regulator's compliance verification requires the office's monitoring intelligence, and they are institutionally distinct because the analytical independence to monitor and the regulatory authority to enforce compliance are different integrity requirements that should not be combined in a single office.

The interface is governed by the intelligence-regulation boundary. The regulator's function requires the regulatory authority, the statutory process, and the power to enforce compliance: the regulator issues the safety standards, the reliability standards, and the environmental standards; the regulator inspects the regulated entities' compliance with the standards; and the regulator enforces the compliance through the fines, the compliance orders, and the licence revocations. The geospatial office's function requires the analytical independence, the methodological rigour, and the freedom to produce intelligence that may contradict the regulator's assumptions or the regulated entities' claims. When the office assumes the regulatory function, it exceeds its mandate and compromises its analytical independence, because the office that is also regulating is tempted to produce intelligence that supports its regulatory preferences. When the regulator assumes the intelligence function, it exceeds its mandate and compromises its regulatory authority, because the regulator that is also producing intelligence is tempted to interpret the intelligence in favour of its regulatory priorities. The boundary is the structural condition that allows both offices to perform their functions with the integrity their functions require.

The interface has two variants, one for each sector the office serves. The energy regulator variant is defined by the asset registry that the regulator provides: the regulator maintains the authoritative record of the jurisdiction's energy infrastructure (the generation facilities, the transmission corridors, the substations, and the pipelines), and the office monitors the infrastructure against the registry, producing the infrastructure risk map, the encroachment alerts, and the maintenance advisories. The minerals regulator variant is defined by the commercial supply relationship: the regulator regulates the mineral sector (the exploration licences, the mining permits, the environmental compliance), and the office is a commercial supplier of mineral intelligence to the sector, selling the prospectivity maps and the resource estimates to the mining and exploration firms. The two variants share the intelligence-regulation boundary, but they differ in the commercial character: the energy variant is a public-service intelligence provision (the office produces the intelligence under the public-service appropriation), and the minerals variant is a commercial intelligence supply (the office produces the intelligence under the commercial charge code).

The interface's institutional location varies across jurisdictions. In some jurisdictions, the geospatial office and the sectoral regulator sit within the same ministry (typically the energy ministry or the mines ministry); in others, the two offices sit in different ministries or different levels of

government. The interface may also span the public-private boundary, with the regulator as a public authority and the office as a public-sector agency that sells commercial intelligence to private-sector customers. The institutional geometry does not change the functional division: the office produces the intelligence, and the regulator regulates the sector, regardless of whether the two offices sit in the same ministry or in different ministries. The institutional geometry does, however, affect the interface's complexity, with the same-ministry variant simplifying the coordination but introducing a risk that the analytical independence is compromised by the ministry's regulatory priorities, and the different-ministry variant preserving the analytical independence but requiring more formal coordination mechanisms.

26.2 The Scope

The interface covers six activities: asset registry sharing, monitoring reporting, compliance verification support, commercial supply governance, framework maintenance, and post-incident review. The activities are differentiated by sector, with the energy regulator variant covering the first, second, third, fifth, and sixth activities, and the minerals regulator variant covering the fourth, fifth, and sixth activities.

Asset registry sharing is the energy regulator variant's foundational activity. The regulator maintains the asset registry as the authoritative record of the jurisdiction's energy infrastructure, and the office consumes the registry as the base layer for its infrastructure monitoring (treated in Chapter 21, section 21.2). The sharing is conducted through the inter-agency framework, which establishes the terms of the registry sharing, the update frequency, and the data format. The regulator provides the registry in the standard format (typically a Shapefile or a GeoPackage with the asset identifiers, the asset types, the geolocations, the criticality ratings, and the attribute data), and the office ingests the registry into its infrastructure asset registry database. The sharing is conducted under the Security Confidential classification, because the registry contains the geolocation of the critical energy infrastructure, and the sharing is restricted to the security-cleared personnel in both offices.

Monitoring reporting is the energy regulator variant's ongoing activity. The office reports its energy infrastructure monitoring results to the regulator and the energy sector entity on the monthly cadence, with the event-driven reporting for the significant encroachments. The reporting comprises the infrastructure risk map (the monthly standing product), the encroachment alerts (the event-driven product), and the maintenance advisories (the quarterly analytical product), each treated in Chapter 21, section 21.4. The reporting is classified Security Confidential, and it is delivered to the regulator and the energy sector entity through the secure distribution channel. The reporting is the operational expression of the office's monitoring function, and the reporting is treated in detail in section 26.4 below.

Compliance verification support is the energy regulator variant's regulatory support activity. The regulator uses the office's monitoring intelligence to verify the energy sector entity's compliance with the safety and reliability standards. The regulator's compliance verification checks whether the entity is maintaining the vegetation clearance zones along the transmission corridors (treated in Chapter 21, section 21.5), whether the entity is responding to the encroachment alerts within the required timeframe (typically 30 days for the high-priority encroachments), and whether the entity's maintenance programme is adequate to prevent the infrastructure's degradation. The office's

monitoring intelligence is the evidence base for the regulator's compliance verification, and the office's intelligence may be used in the regulatory enforcement actions (the fines, the compliance orders, the licence revocations). The office does not participate in the regulatory enforcement: the office produces the intelligence, the regulator enforces the compliance, and the two functions are coordinated but not merged.

Commercial supply governance is the minerals regulator variant's central activity. The office sells the prospectivity maps and the resource estimates (treated in Chapter 22) to the mining and exploration firms in the sector that the regulator regulates. The commercial supply is governed by the commercial firewall (treated in Chapter 5), which separates the commercial arm from the public-science arm and prevents the commercial arm's intelligence from being compromised by the regulator's relationship with the regulated sector. The commercial supply governance comprises the commercial charge code management (every commercial mineral intelligence product is produced under a commercial charge code, with the cost recovered from the customer), the commercial release procedure (the office manager authorizes the release, with the technical lead's validation and the production operator's classification marking), and the commercial customer management (the customer contracts, the licence terms, the confidentiality agreements). The regulator does not have access to the commercial intelligence unless the office manager authorizes the release, and the office manager authorizes the release only under the conditions specified in the commercial firewall (treated in Chapter 5, section 5.4).

Framework maintenance is the standing maintenance of the inter-agency framework between the reporting cycles. The framework comprises the data-sharing agreement (for the energy variant, the asset registry sharing terms; for the minerals variant, the licence registry sharing terms), the reporting protocol (for the energy variant, the monitoring reporting format, cadence, and classification), the commercial supply protocol (for the minerals variant, the commercial charge code management, the release procedure, the customer management), and the review cycle (the annual review of the framework's effectiveness). The framework is maintained by the office manager, who is the office's point of contact for the regulator, and the framework is reviewed annually, after every major incident (for the energy variant), and after every commercial supply dispute (for the minerals variant).

Post-incident review is the energy regulator variant's event-driven review activity, conducted after an energy infrastructure incident (a transmission line failure, a pipeline leak, a substation fire). The review is conducted by the office and the regulator jointly, with each office presenting its perspective on the incident's timeline, the office's monitoring intelligence (whether the monitoring detected the conditions that led to the incident), the regulator's compliance verification (whether the regulator's verification identified the entity's compliance failures), and the lessons learned. The review is documented, and the documentation is the basis for the improvement of the monitoring, the reporting, and the framework for the next incident. The post-incident review is the energy regulator variant's equivalent of the post-event review in the emergency management authority interface (treated in Chapter 24, section 24.8).

The scope does not include the regulatory enforcement (the fines, the compliance orders, the licence revocations), which is the regulator's function. It does not include the infrastructure operation (the generation, the transmission, the distribution of energy), which is the energy sector entity's

function. It does not include the mineral extraction (the exploration, the mining, the processing), which is the mining and exploration firms' function. It does not include the environmental compliance assessment (the environmental impact assessments, the environmental monitoring), which is the environmental regulator's function. The office that expands its scope to include the enforcement, the operation, the extraction, or the environmental compliance exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

26.3 The Method

The method of the interface is the structured exchange of intelligence and regulatory information between the office and the regulator, governed by the inter-agency framework and executed through the asset registry sharing, the monitoring reporting, and the commercial supply governance. The method comprises the asset registry synchronization, the monitoring reporting procedure, the compliance verification support procedure, the commercial supply procedure, and the coordination procedures.

The asset registry synchronization is the energy regulator variant's foundational procedure. The regulator maintains the asset registry as the authoritative record, and the office's copy must be synchronized with the regulator's master. The synchronization is conducted weekly, with the regulator providing the incremental updates (the new assets, the modified assets, the decommissioned assets) through the secure data interface. The office's data pipeline operator retrieves the incremental updates and ingests them into the office's infrastructure asset registry database, with the synchronization logged in the office's data ingestion log. The synchronization includes the conflict resolution: if the office's monitoring has detected a discrepancy between the registry and the actual infrastructure (a decommissioned substation that the registry still shows as active, a new transmission line that the registry does not yet show), the office reports the discrepancy to the regulator through the discrepancy reporting procedure (analogous to the cadastral discrepancy reporting treated in Chapter 23, section 23.6), and the regulator investigates and resolves the discrepancy. The synchronization is monitored for the completeness and the timeliness: if the incremental update is not received by the agreed time (typically Monday 09:00), the operator notifies the office manager, who contacts the regulator's data liaison to resolve the delay.

The monitoring reporting procedure is the energy regulator variant's ongoing procedure. The office produces the infrastructure risk map, the encroachment alerts, and the maintenance advisories on the monthly cadence (treated in Chapter 21, section 21.6), and the products are delivered to the regulator and the energy sector entity through the secure distribution channel. The reporting procedure specifies the product formats (GeoTIFF for the raster products, Shapefile for the vector products, PDF for the report products), the delivery channel (the secure data interface with the mutual TLS authentication), the classification marking (Security Confidential, applied by the production operator), and the delivery timeline (the fifth business day of each month for the standing products, the event-driven timeline for the encroachment alerts). The reporting procedure also specifies the regulator's response: the regulator acknowledges the receipt within 24 hours, the regulator reviews the products within 7 days, and the regulator communicates the compliance verification findings to the office within 30 days. The regulator's response is logged in the office's reporting log, and the response is the input to the next reporting cycle's planning.

The compliance verification support procedure is the energy regulator variant's regulatory support procedure. The regulator uses the office's monitoring intelligence to verify the energy sector entity's compliance with the safety and reliability standards, and the office supports the verification by providing the additional intelligence, the analytical interpretation, and the expert testimony. The additional intelligence may include the high-resolution imagery of the specific infrastructure that the regulator is investigating, the historical monitoring data for the specific time period that the regulator is auditing, or the analytical products (the change detection results, the vulnerability assessments) for the specific assets that the regulator is examining. The analytical interpretation may include the specialist's explanation of the monitoring results, the specialist's assessment of the encroachment's severity, or the specialist's recommendation for the regulatory action. The expert testimony may include the specialist's appearance before the regulatory hearing, the specialist's written statement for the regulatory proceeding, or the specialist's deposition for the enforcement litigation. The office's support is conducted under the intelligence-regulation boundary: the office provides the intelligence, the interpretation, and the testimony, but the office does not participate in the regulatory decision, and the regulator's decision is the regulator's responsibility.

The commercial supply procedure is the minerals regulator variant's central procedure. The office sells the prospectivity maps and the resource estimates to the mining and exploration firms, and the supply is governed by the commercial firewall (treated in Chapter 5). The procedure comprises the customer inquiry reception (the firm's request for the prospectivity map or the resource estimate, with the area of interest, the commodity, and the required deliverables), the commercial charge code establishment (the office manager's authorization of the commercial charge code, with the cost estimate and the customer's agreement), the production (the resource economics specialist's production of the prospectivity map or the resource estimate under the commercial charge code, with the technical lead's validation), the release (the office manager's authorization of the release, with the production operator's classification marking and the licence terms' application), and the delivery (the production operator's delivery of the product to the customer through the commercial distribution channel). The procedure is documented in the office's commercial supply log, with the log retained for 10 years after the completion of the commercial supply (the standard retention period for the commercial minerals commissions, aligned to the contractual disputes' statute of limitations).

The coordination procedures govern the routine and the event-driven coordination between the office and the regulator. The routine coordination is conducted through the periodic coordination meetings (typically quarterly, with the office manager and the regulator's compliance director for the energy variant; annually, with the office manager and the regulator's director for the minerals variant), the joint inspections (for the energy variant, the joint inspections of the critical infrastructure that the office monitors and the regulator regulates), and the informal communication (the office manager's periodic check-ins with the regulator's compliance director). The event-driven coordination is conducted through the incident notification (the office's notification of the regulator when the monitoring detects a significant encroachment or a potential failure), the joint investigation (the joint investigation of the incident, with the office providing the intelligence and the regulator providing the regulatory authority), and the enforcement support (the office's support of the regulator's enforcement action, with the intelligence and the testimony). The coordination procedures specify the contact persons (the office manager is the office's primary point of contact; the regulator's

compliance director is the regulator's primary point of contact), the contact channels, and the response times (the office manager responds to the regulator's notification request within 30 minutes during the business hours and within 2 hours outside the business hours).

26.4 The Products

The interface's products are the office's energy infrastructure and minerals intelligence products, delivered to the regulator and the energy sector entity (for the energy variant) or to the commercial minerals customers (for the minerals variant). The products comprise five types, differentiated by the sector and the consumer.

The infrastructure risk map is the energy regulator variant's standing product, produced on the monthly cadence. The map shows the monitored energy infrastructure (the generation facilities, the transmission corridors, the substations, and the pipelines, overlaid on the high-resolution satellite imagery) and the vulnerability scores (colour-coded by priority: red for high-priority, orange for moderate-priority, yellow for low-priority). The map is used by the regulator for the compliance verification (the regulator's verification that the entity is maintaining the infrastructure in accordance with the safety and reliability standards) and by the energy sector entity for the infrastructure management (the entity's maintenance planning, the resource allocation, the risk assessment). The map carries the Security Confidential classification, because it identifies the critical energy infrastructure's geolocation and its vulnerabilities.

The encroachment alerts are the energy regulator variant's event-driven product, produced as the encroachments are detected. The alerts identify the specific encroachment (the type, the location, the proximity to the asset, the vulnerability score) and provide the supporting imagery (the before-and-after image pair, annotated to show the encroachment). The alerts are used by the regulator for the regulatory enforcement (the regulator's verification that the entity is responding to the encroachments within the required timeframe) and by the energy sector entity for the encroachment response (the entity's dispatch of the maintenance crews to the encroachment sites). The alerts carry the Security Confidential classification.

The maintenance advisories are the energy regulator variant's analytical product, produced for the energy sector entity's maintenance planning and the regulator's compliance assessment. The advisories identify the infrastructure at risk from the ground conditions (the subsidence, the erosion, the slope instability) and recommend the maintenance actions (the vegetation clearing, the ground stabilization, the structural reinforcement). The advisories are produced quarterly, synthesizing the monthly monitoring results and the ground condition analysis into a longer-term planning document. The advisories carry the Security Confidential classification.

The compliance verification report is the energy regulator variant's regulatory support product, produced at the regulator's request for the specific compliance verification or the enforcement action. The report compiles the office's monitoring intelligence for the specific infrastructure, the specific time period, or the specific encroachment that the regulator is investigating, with the analytical interpretation and the expert assessment. The report may include the high-resolution imagery, the change detection results, the vulnerability assessments, the historical monitoring data, and the specialist's written assessment. The report carries the Security Confidential classification, and the

report is delivered to the regulator through the secure distribution channel, with the access restricted to the regulator's authorized personnel and the office's authorized personnel.

The prospectivity map and the resource estimate are the minerals regulator variant's commercial products, produced for the commercial minerals customers. The prospectivity map identifies the areas with the high mineral potential, based on the geological, the geochemical, and the geophysical data, and the resource estimate quantifies the mineral resource in the identified areas, based on the drilling data and the geostatistical analysis. The products are produced under the commercial charge code, with the cost recovered from the customer, and the products carry the Confidential or the Security Confidential classification depending on the customer's request and the product's content. The products are delivered to the customer through the commercial distribution channel, with the licence terms' application. The regulator does not have access to the commercial products unless the office manager authorizes the release under the conditions specified in the commercial firewall.

26.5 The Consumers

The interface's primary consumers are the sectoral regulators themselves. The energy regulator uses the office's monitoring intelligence for the compliance verification, the regulatory enforcement, and the performance assessment. The regulator's compliance verification division uses the infrastructure risk map, the encroachment alerts, and the compliance verification reports to verify the energy sector entity's compliance with the safety and reliability standards. The regulator's enforcement division uses the encroachment alerts and the compliance verification reports to support the enforcement actions (the fines, the compliance orders, the licence revocations). The regulator's performance assessment division uses the infrastructure risk map and the maintenance advisories to assess the entity's overall performance and the regulator's regulatory effectiveness.

The minerals regulator uses the office's public-science minerals intelligence (the geological maps, the mineral resource assessments, the regional prospectivity studies) for the regulatory policy, the licence management, and the sectoral planning. The regulator's policy division uses the public-science intelligence to develop the mineral sector's regulatory framework (the licensing rules, the environmental standards, the royalty rates). The regulator's licence management division uses the public-science intelligence to evaluate the licence applications and the licence renewals. The regulator's sectoral planning division uses the public-science intelligence to plan the sector's development (the exploration targets, the mining zones, the infrastructure requirements). The regulator does not use the office's commercial minerals intelligence, because the commercial intelligence is produced under the commercial firewall and is not accessible to the regulator unless the office manager authorizes the release.

The energy sector entity is the secondary consumer of the energy regulator variant. The entity uses the office's monitoring intelligence for the infrastructure management, the encroachment response, and the maintenance planning, as treated in Chapter 21, section 21.5. The entity receives the intelligence through the regulator's distribution (the regulator distributes the intelligence to the entity as part of the regulator's compliance verification), and the entity does not receive the intelligence directly from the office unless the inter-agency framework specifies the direct distribution. The entity's use of the intelligence is governed by the regulator's compliance framework: the entity must respond to the encroachment alerts within the required timeframe, the entity must maintain the

infrastructure in accordance with the standards, and the entity's failure to do so is the regulatory violation that the regulator enforces.

The commercial minerals customers are the secondary consumers of the minerals regulator variant. The customers are the mining and exploration firms that purchase the office's prospectivity maps and resource estimates, and the customers use the intelligence for the exploration decisions, the mining investments, and the resource declarations. The customers receive the intelligence through the commercial distribution channel, and the customers' use of the intelligence is governed by the licence terms and the confidentiality agreements. The customers' relationship with the office is a commercial relationship, not a regulatory relationship, and the regulator is not a party to the commercial relationship.

The public is the indirect beneficiary of the interface's intelligence, but the public is not the office's direct consumer. The public benefits from the regulator's enforcement of the safety and reliability standards (which the office's intelligence supports) and from the regulator's management of the mineral sector (which the office's public-science intelligence supports). The public receives the intelligence through the regulator's public communication: the regulator's publication of the compliance statistics, the regulator's publication of the enforcement actions, and the regulator's publication of the sectoral performance reports. The office's intelligence underpins the regulator's public communication, but the office does not communicate directly with the public on the regulatory matters, because the public communication on the regulatory matters is the regulator's function.

The political leadership is the tertiary consumer, receiving the office's intelligence through the regulator's briefings. The political leadership (the minister, the senior appointed officials) receives the sectoral briefings from the regulator, and the office's intelligence is summarized in the briefings. The political leadership uses the briefings to make the policy decisions (the regulatory reforms, the budget allocations, the strategic priorities). The office does not brief the political leadership directly on the regulatory matters, because the political leadership's briefings must be consistent with the regulator's regulatory picture, and the regulator is the single point of contact for the political leadership's briefings on the sectoral matters.

26.6 The Cadence

The interface operates on three cadences: the standing cadence, the event-driven cadence, and the review cadence. The cadences are differentiated by the sector, with the energy regulator variant operating on the monthly standing cadence and the event-driven cadence, and the minerals regulator variant operating on the project-based commercial cadence and the annual review cadence.

The standing cadence for the energy regulator variant comprises the monthly monitoring reporting (the office produces the infrastructure risk map, the encroachment alerts, and the maintenance advisories, and delivers them to the regulator and the energy sector entity by the fifth business day of each month), the quarterly compliance verification review (the regulator reviews the office's monitoring intelligence and communicates the compliance verification findings to the office within 30 days of the quarterly reporting cycle), the quarterly coordination meeting (the office manager and the regulator's compliance director review the recent monitoring results, the recent compliance verification findings, the recent enforcement actions, and the framework's operation), the annual asset registry reconciliation (the office and the regulator jointly reconcile the office's copy of

the asset registry with the regulator's master, identifying and resolving the discrepancies), and the annual framework review (the office and the regulator jointly review the framework's effectiveness, identifying the issues encountered and the corrective actions).

The event-driven cadence for the energy regulator variant is activated when an energy infrastructure incident occurs or when the monitoring detects a significant encroachment. The incident notification procedure is activated: the office's hazard specialist notifies the office manager of the incident or the significant encroachment, the office manager notifies the regulator's compliance director within 30 minutes during the business hours and within 2 hours outside the business hours, and the regulator's compliance director acknowledges the notification within 24 hours. The joint investigation procedure may be activated: the office and the regulator jointly investigate the incident, with the office providing the intelligence and the regulator providing the regulatory authority. The enforcement support procedure may be activated: the office supports the regulator's enforcement action, with the intelligence and the testimony. The event-driven cadence is maintained by the hazard specialist, who produces the event-specific intelligence, and the office manager, who coordinates with the regulator.

The commercial cadence for the minerals regulator variant is project-based, with each commercial supply project conducted on its own timeline. The project timeline comprises the customer inquiry reception (the firm's request, with the area of interest, the commodity, and the required deliverables), the commercial charge code establishment (the office manager's authorization, with the cost estimate and the customer's agreement, typically within 5 business days of the inquiry), the production (the resource economics specialist's production, with the timeline determined by the project's scope, typically 4 to 12 weeks for a regional prospectivity map, 8 to 24 weeks for a detailed resource estimate), the release (the office manager's authorization, with the technical lead's validation and the production operator's classification marking, typically within 5 business days of the production completion), and the delivery (the production operator's delivery, with the licence terms' application, typically within 2 business days of the release authorization). The commercial cadence is maintained by the resource economics specialist, who produces the intelligence, and the office manager, who authorizes the charge code and the release.

The review cadence for both variants comprises the annual framework review (the office and the regulator jointly review the framework's effectiveness, identifying the issues encountered and the corrective actions) and the post-incident review (for the energy variant, the joint review of the energy infrastructure incident, treated in section 26.2 above) or the post-supply review (for the minerals variant, the joint review of the commercial supply project, with the office and the customer assessing the product's quality, the timeline's adherence, and the commercial relationship's effectiveness). The review cadence is maintained by the office manager, who schedules and chairs the reviews, and the regulator's compliance director (for the energy variant) or the customer's project manager (for the minerals variant), who participates in the reviews.

The cadence is the mechanism that maintains the interface's rhythm and the interface's improvement over time. The standing cadence maintains the regular intelligence provision. The event-driven cadence delivers the intelligence during the incidents. The commercial cadence delivers the intelligence under the commercial terms. The review cadence improves the interface after the incidents and the commercial projects. The four cadences together constitute the interface's

operational rhythm, and the rhythm is what allows the office to maintain the interface with the regulator across the two sectors, despite the different cadences and the different commercial characters of the two variants.

26.7 The Classification Regime

The interface operates under a tiered classification regime, with the classification determined by the product's content, the product's consumer, and the product's commercial character. The tiered regime comprises three levels: Public, Internal, and Security Confidential, with the additional Confidential level for the commercial minerals intelligence.

The Public classification applies to the public-science minerals intelligence (the geological maps, the mineral resource assessments, the regional prospectivity studies) and the public summaries of the energy regulator's compliance statistics. These products are intended for the public awareness and the sector's general information, and they do not include the operational detail that could compromise the regulatory enforcement or the commercial confidentiality. The Public geological map shows the regional geology, the mineral occurrences, and the exploration targets at the regional scale, without the parcel-level detail or the customer-specific intelligence. The Public compliance statistics show the regulator's enforcement actions at the aggregate level (the number of fines, the number of compliance orders, the number of licence revocations), without the entity-specific detail. The Public classification allows the regulator to publish the products on its website and to use them in the public communication.

The Internal classification applies to the operational versions of the public-science minerals intelligence and the office's internal compliance verification reports. The Internal geological map includes the parcel-level detail and the customer-specific intelligence that the public version excludes. The Internal compliance verification report compiles the office's monitoring intelligence for the regulator's internal use, with the analytical interpretation and the expert assessment. The Internal classification restricts the products to the regulator's authorized personnel and the office's authorized personnel, and the products are delivered through the secure distribution channel with the access control.

The Security Confidential classification applies to the energy infrastructure intelligence (the infrastructure risk map, the encroachment alerts, the maintenance advisories, the compliance verification reports) and the asset registry itself. These products include the geolocation of the critical energy infrastructure, which is security-sensitive because the disclosure could inform the attacks on the infrastructure. The Security Confidential classification is consistent with the critical infrastructure information protection frameworks (treated in Chapter 21, section 21.7), and it restricts the products to the security-cleared personnel in the regulator, the energy sector entity, and the office. The classification is maintained by the production operator, who applies the markings, and the office manager, who reviews the markings before the release.

The Confidential classification applies to the commercial minerals intelligence (the customer-specific prospectivity maps and the resource estimates), which is commercial-sensitive because its disclosure could compromise the customer's competitive position. The customer's prospectivity map identifies the areas with the high mineral potential, and the disclosure to the customer's competitors would allow the competitors to acquire the licences in the identified areas before the customer. The

customer's resource estimate quantifies the mineral resource in the customer's licence area, and the disclosure to the customer's competitors or to the market would allow the competitors or the market to anticipate the customer's mining investment decisions. The Confidential classification restricts the products to the customer's authorized personnel and the office's authorized personnel, and the products are delivered through the commercial distribution channel with the access control. The Confidential classification is the commercial integrity guarantee for the commercial minerals intelligence, and the marking is non-negotiable.

The classification regime's tiered structure allows the office to produce the intelligence at the appropriate level of restriction for each audience and each product. The Public products inform the public awareness. The Internal products inform the regulator's operational work. The Security Confidential products inform the critical infrastructure protection. The Confidential products inform the commercial minerals customers. The classification regime is maintained by the production operator, who applies the markings, and the office manager, who reviews the markings before the release. The regime is non-negotiable: every product carries the appropriate classification, and the marking is applied in the standard format (capital letters, top and bottom of every page, in the metadata) and in the standard location. The office manager may declassify a product only if the product's content does not include the operational, the security-sensitive, or the commercial-sensitive information, and the declassification is documented in the production log with the rationale.

26.8 The Comparative Practice

The comparative practice of the sectoral regulator interface spans two sectors (energy and minerals) and multiple jurisdictions, with the institutional arrangements varying by the sector and the jurisdiction. The comparative practice is organized by the sector, with the energy regulator interface examined first and the minerals regulator interface examined second.

The energy regulator interface is exemplified by the United States, the European Union, and Australia. In the United States, the Federal Energy Regulatory Commission (FERC) regulates the interstate transmission of electricity and natural gas, and the North American Electric Reliability Corporation (NERC) develops and enforces the reliability standards for the bulk power system. FERC and NERC use the monitoring intelligence from the Department of Energy's Office of Cybersecurity, Energy Security, and Emergency Response (CESER), the National Geospatial-Intelligence Agency (NGA), and the commercial satellite operators (Maxar, Planet) to verify the energy sector entities' compliance with the safety and reliability standards. The interface's intelligence-regulation boundary is maintained: the monitoring agencies produce the intelligence, FERC and NERC regulate the sector, and the two functions are coordinated through the Department of Energy's sector-specific agency role (treated in Chapter 21, section 21.8). The interface's classification regime is the Security Confidential classification under the Protected Critical Infrastructure Information (PCII) programme.

In the European Union, the Agency for the Cooperation of Energy Regulators (ACER) coordinates the national energy regulators, and the European Network of Transmission System Operators for Electricity (ENTSO-E) develops the network codes. The national energy regulators (such as the Bundesnetzagentur in Germany, the Commission de Régulation de l'Énergie in France, and the Autorità per l'Energia Elettrica e il Gas in Italy) regulate the national energy sectors, and they

use the monitoring intelligence from the European Commission's Joint Research Centre (JRC), the Copernicus Emergency Management Service (CEMS), and the national geospatial agencies to verify the entities' compliance. The interface's intelligence-regulation boundary is maintained: the monitoring agencies produce the intelligence, the national regulators regulate the sector, and the coordination is conducted through the ACER and the ENTSO-E. The interface's classification regime is the restricted information protection under the Directive 2008/114/EC on the European critical infrastructures.

In Australia, the Australian Energy Regulator (AER) regulates the energy markets and the networks, and the Australian Energy Market Operator (AEMO) operates the energy markets and the power system. The AER and the AEMO use the monitoring intelligence from the Geoscience Australia (the national geoscience agency, which produces the satellite-based hazard and infrastructure monitoring), the Commonwealth Scientific and Industrial Research Organisation (CSIRO), and the state geospatial agencies to verify the energy sector entities' compliance. The interface's intelligence-regulation boundary is maintained: the monitoring agencies produce the intelligence, the AER and the AEMO regulate and operate the sector, and the coordination is conducted through the Council of Australian Governments' Energy Council. The interface's classification regime is the restricted information protection under the Australian Security Intelligence Organisation Act 1979 and the Critical Infrastructure Protection framework.

The minerals regulator interface is exemplified by Australia, Canada, and the United States. In Australia, the state mineral resources departments (such as the Geological Survey of New South Wales, the Geological Survey of Queensland, and the Geological Survey of Western Australia) regulate the mineral sector within their states, and Geoscience Australia supplies the mineral intelligence to the sector through its commercial minerals division. The state departments issue the exploration licences and the mining permits, regulate the environmental compliance, and collect the royalties. Geoscience Australia produces the public-science minerals intelligence (the geological maps, the mineral resource assessments, the regional prospectivity studies) under the public-service appropriation, and the commercial minerals intelligence (the customer-specific prospectivity maps and the resource estimates) under the commercial charge code. The interface's intelligence-regulation boundary is maintained: the state departments regulate the sector, Geoscience Australia supplies the intelligence, and the commercial firewall prevents the commercial intelligence from being compromised by the regulatory relationship. The interface's classification regime is the Public classification for the public-science intelligence and the Confidential classification for the commercial intelligence.

In Canada, the provincial mineral resources departments (such as the Ontario Geological Survey, the Ministère de l'Énergie et des Ressources Naturelles du Québec, and the British Columbia Geological Survey) regulate the mineral sector within their provinces, and the Geological Survey of Canada (GSC, within Natural Resources Canada) supplies the mineral intelligence to the sector. The provincial departments issue the exploration licences and the mining permits, regulate the environmental compliance, and collect the royalties. The GSC produces the public-science minerals intelligence under the public-service appropriation, and the provincial geological surveys produce the regional intelligence under their own appropriations. The commercial minerals intelligence is supplied by the commercial geological consultancies (such as the S&P Global Market Intelligence,

the Wood Mackenzie, and the MinEx Consulting), not by the GSC, and the commercial supply is not subject to the government commercial firewall. The Canadian model is a variant of the Australian model, with the commercial supply shifted from the government agency to the private sector.

In the United States, the Bureau of Land Management (BLM, within the Department of the Interior) regulates the mineral sector on the federal lands (which comprise approximately 28 per cent of the U.S. land area, primarily in the western states), and the United States Geological Survey (USGS, also within the Department of the Interior) supplies the mineral intelligence to the sector. The BLM issues the exploration licences and the mining permits, regulates the environmental compliance, and collects the royalties. The USGS produces the public-science minerals intelligence under the public-service appropriation, and the USGS does not produce the commercial minerals intelligence (the commercial intelligence is supplied by the private sector). The U.S. model is similar to the Canadian model, with the regulator (BLM) and the intelligence producer (USGS) both within the Department of the Interior but with distinct functions. The U.S. model is the example of the integrated variant (the regulator and the intelligence producer within the same department), and the model's intelligence-regulation boundary is maintained through the functional separation rather than the institutional separation.

The comparative practice across the sectors and the jurisdictions shows that the intelligence-regulation boundary is the shared element of the sectoral regulator interface. The boundary is maintained through the institutional separation (the Australian and the Canadian models), the functional separation (the U.S. model), or the hybrid arrangements (the European Union's ACER-ENTSO-E-JRC arrangement). The boundary is the condition of both the office's analytical independence and the regulator's regulatory authority, and the boundary is the subject of this chapter. The geospatial office operates under whichever model the jurisdiction has adopted, and the office's interface with the regulator is shaped by the model. The model determines the institutional complexity, the commercial character, and the classification regime, but the intelligence-regulation boundary is the same across the models.

26.9 Worked Example: The Highmoor-Portsmith Transmission Corridor Compliance Verification

The compliance verification is conducted for the Highmoor-Portsmith transmission corridor, the 220-kilometre, 500-kV transmission line treated in Chapter 21, section 21.9. The corridor's monitoring in July 2024 detected 12 changes within the right-of-way, of which 11 were confirmed (7 vegetation encroachments, 3 third-party constructions, 1 ground subsidence) and 1 was excluded as a false alarm. The encroachment alerts were delivered to the energy sector entity (the Highmoor Power Company) and the sectoral regulator (the Energy Regulatory Commission) on 5 July 2024, with 8 high-priority alerts (4 vegetation encroachments within the 15-metre clearance zone, 1 ground subsidence near a critical tower, and 3 third-party constructions within the right-of-way).

The regulator's compliance verification begins on 8 July 2024, three days after the encroachment alerts' delivery. The regulator's compliance director reviews the alerts and identifies the 8 high-priority encroachments as the priority cases for the compliance verification. The regulator's compliance officer is assigned to the verification, and the officer requests the office's compliance verification report for the 8 high-priority encroachments through the secure distribution channel. The

request specifies the encroachments' identifiers (the office's alert numbers), the required deliverables (the high-resolution imagery, the change detection results, the vulnerability assessments, the historical monitoring data for the previous 12 months, and the specialist's written assessment), and the deadline (15 July 2024, 7 days from the request).

The office's hazard specialist produces the compliance verification report on 10-14 July 2024. The report compiles the monitoring intelligence for the 8 high-priority encroachments, with each encroachment's section comprising the encroachment's identifier, the encroachment's type (vegetation encroachment, third-party construction, or ground subsidence), the encroachment's location (the latitude, the longitude, the nearest tower identifier, the distance to the tower), the encroachment's detection date and the detection method (the object-based change detection on the Maxar WorldView-3 imagery, treated in Chapter 21, section 21.9), the encroachment's vulnerability score (the Criticality Weight × Proximity Factor × Change Type Factor, treated in Chapter 21, section 21.3), the before-and-after image pair (annotated to show the encroachment), the historical monitoring data (the previous 12 months' change detection results for the same location, showing no prior encroachment), and the specialist's written assessment (the encroachment's severity, the recommended action, and the regulatory implication). The report's summary section identifies the overall compliance picture: the 8 high-priority encroachments represent a significant compliance concern, with the 4 vegetation encroachments within the clearance zone being the regulatory violations (the safety standard requires the 15-metre clearance, and the encroachments are within the clearance), the 1 ground subsidence being a potential geotechnical concern (the subsidence may affect the tower's foundation, and the standard requires the geotechnical investigation), and the 3 third-party constructions being the right-of-way violations (the right-of-way is the regulated zone, and the constructions within the right-of-way require the regulator's authorization, which was not obtained).

The production operator produces the report in the agreed format (PDF for the report, GeoTIFF for the imagery, Shapefile for the encroachment locations), applies the Security Confidential classification marking, and submits the draft to the office manager. The office manager reviews the report on 14 July 2024, verifying the classification marking, the analytical content, and the specialist's assessment. The office manager authorizes the release, and the production operator delivers the report to the regulator's compliance officer through the secure distribution channel on 15 July 2024, meeting the deadline.

The regulator's compliance verification proceeds on 16-25 July 2024. The compliance officer reviews the report and conducts the verification. The verification checks whether the Highmoor Power Company (the energy sector entity) has responded to the encroachment alerts within the required timeframe (the standard requires the response within 30 days of the alert, and the alerts were delivered on 5 July 2024, so the response is due by 4 August 2024). The verification also checks whether the entity's response is adequate: the 4 vegetation encroachments within the clearance zone require the vegetation clearing (the standard requires the clearing within 30 days of the alert), the 1 ground subsidence requires the geotechnical investigation (the standard requires the investigation within 60 days of the alert), and the 3 third-party constructions require the legal action (the standard requires the entity to notify the regulator and to initiate the legal proceedings within 30 days of the alert).

The compliance officer conducts the field inspection on 20 July 2024, visiting the 8 high-priority encroachment sites with the entity's representatives. The field inspection confirms the office's monitoring intelligence: the 4 vegetation encroachments are within the clearance zone (the office's measurement of the vegetation height and the distance to the transmission line is confirmed), the 1 ground subsidence is visible (the office's measurement of the subsidence magnitude is confirmed), and the 3 third-party constructions are within the right-of-way (the office's identification of the constructions' locations is confirmed). The field inspection's findings are documented in the regulator's inspection report, and the office's monitoring intelligence is cited as the evidence base for the findings.

The regulator's enforcement decision is made on 25 July 2024. The compliance officer recommends the enforcement action to the regulator's enforcement division, based on the office's monitoring intelligence and the field inspection's findings. The enforcement division reviews the recommendation and issues the compliance order on 30 July 2024: the Highmoor Power Company is ordered to clear the 4 vegetation encroachments within 30 days (by 29 August 2024), to conduct the geotechnical investigation for the 1 ground subsidence within 60 days (by 28 September 2024), and to initiate the legal proceedings for the 3 third-party constructions within 30 days (by 29 August 2024). The compliance order also includes the fine of $250,000 for the violations (the standard fine for the high-priority encroachment violations, calculated as $50,000 per vegetation encroachment within the clearance zone and $25,000 per third-party construction within the right-of-way).

The regulator's enforcement decision is communicated to the office on 1 August 2024, as the decision feedback. The feedback comprises the decision (the compliance order and the fine), the rationale (the violations of the safety standard and the right-of-way regulation, based on the office's monitoring intelligence and the field inspection), the implementation (the compliance order's deadlines and the fine's payment schedule), and the outcome (to be communicated after the compliance order's deadlines). The office's hazard specialist reviews the feedback and notes that the office's monitoring intelligence was used as the evidence base for the enforcement action, which confirms the intelligence's usefulness for the regulator's compliance verification. The specialist also notes that the office's monitoring intelligence was accurate (the field inspection confirmed the office's measurements), which validates the monitoring method's reliability.

The compliance order's implementation is tracked through the regulator's monthly compliance reporting. The Highmoor Power Company clears the 4 vegetation encroachments by 25 August 2024 (4 days before the deadline), conducts the geotechnical investigation by 20 September 2024 (8 days before the deadline), and initiates the legal proceedings for the 3 third-party constructions by 28 August 2024 (1 day before the deadline). The fine is paid on 15 August 2024. The implementation's outcome is communicated to the office on 1 October 2024: the compliance order's requirements were met within the deadlines, the violations were resolved, and the regulator's enforcement action was effective. The office's monitoring intelligence was the evidence base for the entire enforcement process, from the initial encroachment detection to the final compliance verification.

The post-incident review is conducted on 15 November 2024, jointly by the office and the regulator. The review covers the timeline (the encroachment detection on 3 July 2024, the alert delivery on 5 July, the compliance verification request on 8 July, the report delivery on 15 July, the field inspection on 20 July, the enforcement decision on 25 July, the compliance order on 30 July, the

implementation by 28 September, the fine payment on 15 August), the products delivered (the encroachment alerts and the compliance verification report, both Security Confidential), the issues encountered (the 30-day response timeframe was adequate for the vegetation clearing but tight for the legal proceedings, and the regulator's enforcement division suggested the longer timeframe for the legal proceedings in the future), the lessons learned (the office's monitoring intelligence was accurate and useful, the field inspection confirmed the intelligence, the enforcement was effective, and the longer timeframe for the legal proceedings should be considered in the next revision of the safety standard), and the corrective actions (the office's monitoring will continue on the monthly cadence, the regulator's compliance verification will be conducted within 7 days of the alert delivery, and the safety standard's revision will be considered in the next annual review). The review is documented, and the documentation is the basis for the improvement of the interface, the products, and the procedures for the next compliance verification.

This worked example illustrates the interface's full operation: the encroachment detection by the office's monitoring (the 12 changes, the 11 confirmed, the 8 high-priority), the alert delivery to the regulator and the entity (5 July 2024, Security Confidential), the regulator's compliance verification request (8 July), the office's compliance verification report production (10-14 July, with the 4-section structure per encroachment and the summary section), the report delivery (15 July, meeting the deadline), the regulator's compliance verification (16-25 July, with the field inspection on 20 July confirming the office's intelligence), the enforcement decision (25 July), the compliance order (30 July, with the deadlines and the fine), the implementation (by 28 September), the outcome communication (1 October), and the post-incident review (15 November). The interface's intelligence-regulation boundary was maintained throughout: the office produced the intelligence, the regulator regulated the sector, and the two functions were coordinated through the secure distribution channel, the compliance verification report, and the post-incident review. The interface's classification regime was maintained: the Security Confidential marking on the alerts and the report, the restricted distribution, and the access control. The interface's compliance verification support was effective: the office's monitoring intelligence was the evidence base for the regulator's enforcement action, and the enforcement action resolved the violations and ensured the corridor's safety. The interface is the foundation of the office's energy infrastructure intelligence delivery to the sectoral regulator, and the interface's integrity is the condition of the office's contribution to the energy sector's regulation. The Highmoor-Portsmith corridor compliance verification demonstrates the interface's value: the office's monitoring intelligence, produced on the monthly cadence with the high-resolution satellite imagery and the object-based change detection, detected the encroachments that the entity had not reported, the regulator's enforcement action required the encroachments' resolution, and the corridor's safety was maintained. The 2003 Northeast Blackout, triggered by a transmission line sagging into overgrown vegetation in Ohio (treated in Chapter 21, section 21.1), demonstrated the catastrophic consequences of the undetected encroachment: 55 million people without power, $6 billion in economic damage, and a cascading failure that took four days to fully restore. The geospatial office's energy infrastructure intelligence, delivered to the sectoral regulator through the interface treated in this chapter, is the mechanism by which such encroachments are detected, reported, and resolved before they cause failures, and the detection and the resolution are the interface's contribution to the energy system's reliability and the public's safety.

Chapter 27 — The Water Authority Interface

27.1 The Interface Defined

The water authority interface is the boundary between the office that produces flood risk models and the authority that governs the waterways and the waterway infrastructure. The geospatial office produces the flood hazard maps, the flood inundation maps, and the flood exposure reports treated in Chapter 19, and the water authority manages the waterways that convey the floodwaters: the rivers, the canals, the channels, the levees, the floodgates, the pumping stations, the dams, and the reservoirs. The two functions are operationally interdependent because the office's flood models require the water authority's waterway data (the channel geometry, the levee elevations, the gate and pump operations, the reservoir releases), and the water authority's waterway management requires the office's flood predictions (the forecast river stages, the forecast inundation extents, the forecast flood durations). The two functions are institutionally distinct because the analytical independence to model the flood and the operational authority to manage the waterway are different integrity requirements that should not be combined in a single office.

The interface is governed by the flood-water boundary. The water authority's function requires the operational authority over the waterways: the authority maintains the levees (the inspections, the repairs, the upgrades), operates the floodgates and the pumping stations (the opening and the closing, the pumping schedules, the emergency operations), manages the water allocations (the irrigation, the municipal supply, the environmental flows), and operates the dams and the reservoirs (the release decisions, the storage management, the flood control operations). The geospatial office's function requires the analytical independence, the methodological rigour, and the freedom to produce flood models that may contradict the water authority's operational assumptions: the office's model may predict the flood stage that exceeds the levee's design elevation, the office's model may identify the levee's vulnerability that the authority's inspection has missed, or the office's model may recommend the evacuation that the authority's operational plan does not anticipate. When the office assumes the operational function, it exceeds its mandate and compromises its analytical independence, because the office that is also operating the waterways is tempted to produce models that support its operational preferences. When the authority assumes the intelligence function, it exceeds its mandate and compromises its operational authority, because the authority that is also producing the flood models is tempted to interpret the models in favour of its operational convenience. The boundary is the structural condition that allows both offices to perform their functions with the integrity their functions require.

The interface's institutional location varies across jurisdictions. In some jurisdictions, the geospatial office and the water authority sit within the same ministry (typically the environment ministry, the water ministry, or the public works ministry); in others, the two offices sit in different ministries or different levels of government. The interface may also span the national-subnational boundary, with the office at the national level and the water authority at the sub-national level (a common arrangement in the federal jurisdictions, where the water management is a state or provincial responsibility), or vice versa. The institutional geometry does not change the functional division: the office produces the flood models, and the authority manages the waterways, regardless of whether the two offices sit in the same ministry or in different ministries. The institutional geometry does,

however, affect the interface's complexity, with the same-ministry variant simplifying the coordination but introducing a risk that the analytical independence is compromised by the ministry's operational priorities, and the different-ministry variant preserving the analytical independence but requiring more formal coordination mechanisms.

The interface is bilateral in its primary geometry: the geospatial office and the water authority are the two parties. The interface is multilateral in its operational expression, because the water authority's waterway management affects and is affected by the emergency management authority (treated in Chapter 24, with the authority's waterway operations supporting the emergency response), the sectoral regulator (treated in Chapter 26, with the regulator regulating the water sector's safety and reliability), the national statistics agency (treated in Chapter 28, with the agency's data informing the water authority's water allocation decisions), and the cadastral authority (treated in Chapter 23, with the cadastral data informing the water authority's levee and channel management). The office's flood models are delivered to the water authority, and the water authority coordinates the waterway management with the other agencies. The office does not coordinate directly with the other agencies on the waterway management, because the coordination is the water authority's function, and the office's role is to produce the flood models that inform the authority's coordination.

27.2 The Scope

The interface covers six activities: waterway data sharing, flood model integration, event coordination, levee and infrastructure monitoring, framework maintenance, and post-event review. The activities are differentiated by the data flow's direction, the activity's cadence, and the activity's event-driven character.

Waterway data sharing is the water-authority-to-office data flow that provides the office with the waterway data its flood models require. The water authority provides the channel geometry (the cross-sections, the longitudinal profiles, the roughness coefficients, surveyed by the authority's hydrographic teams and updated as the channels are maintained), the levee data (the levee alignments, the crest elevations, the design standards, the inspection results, maintained by the authority's levee programme), the hydraulic structure data (the floodgates, the pumping stations, the weirs, the dams, with the operating curves, the discharge ratings, and the operating rules), the reservoir data (the storage capacities, the current storage levels, the release schedules, the flood control rules), and the water allocation data (the allocations, the diversions, the returns, the consumptive use). The office consumes the waterway data for its hydrological and hydraulic models (the HEC-HMS and the HEC-RAS models treated in Chapter 10, section 10.2, and Chapter 19, section 19.3), with the channel geometry as the hydraulic model's geometric input, the levee data as the model's levee-failure input, the hydraulic structure data as the model's boundary condition, the reservoir data as the model's upstream boundary condition, and the water allocation data as the model's baseflow input. The waterway data sharing is conducted on the standing cadence (annually for the channel geometry and the levee data, monthly for the reservoir storage, daily for the reservoir releases and the water allocations during the flood season).

Flood model integration is the bidirectional activity that integrates the office's flood models with the water authority's waterway management models. The office's flood models simulate the rainfall-runoff process (the HEC-HMS model, producing the river flows from the rainfall) and the inundation

(the HEC-RAS model, producing the inundation extents from the river flows). The water authority's models simulate the waterway hydraulics (the authority's own hydraulic models, which may be the HEC-RAS models or the proprietary models such as the Delft3D or the MIKE 11) and the gate and pump operations (the authority's operational models, which simulate the gates' and the pumps' response to the river stages). The integration is conducted through the data interfaces: the office's HEC-RAS model consumes the water authority's waterway data, and the water authority's operational models consume the office's flood predictions. The integration is the mechanism that allows the two offices to produce a coherent flood risk picture for the jurisdiction, despite the functional boundary between them.

Event coordination is the event-driven activity that coordinates the office's flood modelling with the water authority's waterway operations during a flood event. The coordination comprises the communication channels (the office manager and the authority's operations director, with the contact persons, the contact channels, and the response times specified in the inter-agency framework), the data sharing (the accelerated waterway data sharing, with the authority providing the current reservoir levels, the gate operations, and the levee conditions, and the office providing the current flood models and the forecast inundation extents), the model integration (the accelerated model integration, with the two offices' models exchanging the data in near-real-time), and the escalation procedures (the escalation of the issues that cannot be resolved at the operational level to the two offices' parent authorities). The event coordination is treated in detail in section 27.5 below.

Levee and infrastructure monitoring is the office's satellite-based monitoring of the water authority's waterway infrastructure, complementing the authority's in-situ inspections. The office uses the high-resolution satellite imagery (the Maxar WorldView-3 or the Airbus Pleiades Neo, treated in Chapter 21, section 21.3) and the InSAR analysis (the Sentinel-1 radar data, treated in Chapter 21, section 21.3) to monitor the levees (the vegetation encroachment, the erosion, the subsidence), the floodgates and the pumping stations (the physical changes, the access conditions), and the channels (the bank erosion, the sediment deposition, the vegetation growth). The monitoring is conducted on the monthly cadence, with the event-driven monitoring during the flood events (the rapid post-event monitoring of the levee conditions, the breach detection). The monitoring products are delivered to the water authority as the infrastructure risk maps, the encroachment alerts, and the maintenance advisories, with the products carrying the Internal classification by default and the Security Confidential classification if the products identify the critical infrastructure's vulnerabilities.

Framework maintenance is the standing maintenance of the inter-agency framework between the data exchanges and the event coordinations. The framework comprises the data-sharing agreement (which specifies the waterway data the authority provides, the flood models the office provides, the formats, the delivery channels, and the access controls), the model integration protocol (which specifies the data interfaces, the validation protocols, and the version control), the event coordination protocol (which specifies the communication channels, the contact persons, the response times, and the escalation procedures), the monitoring reporting protocol (which specifies the monitoring products' format, cadence, and classification), and the review cycle (which specifies the annual review of the framework's effectiveness and the post-event review after each major flood event). The framework is maintained by the office manager, who is the office's point of contact for the authority, and the framework is reviewed annually and after every major flood event.

Post-event review is the joint review of the flood event response after the event is resolved. The review is conducted by the office and the authority jointly, with each office presenting its perspective on the event's timeline, the office's flood models (whether the models' predictions matched the observed river stages and the observed inundation extents), the authority's waterway operations (whether the reservoir releases, the gate operations, and the levee performance were effective), the coordination between the two offices (whether the data sharing, the model integration, and the communication were effective), and the lessons learned. The review is documented, and the documentation is the basis for the improvement of the flood models, the waterway operations, and the framework for the next event. The post-event review is the water authority interface's equivalent of the post-event review in the emergency management authority interface (treated in Chapter 24, section 24.8).

The scope does not include the waterway operation (the gate operations, the pump operations, the reservoir releases), which is the water authority's function. It does not include the water allocation (the irrigation, the municipal supply, the environmental flows), which is the water authority's function. It does not include the levee maintenance (the inspections, the repairs, the upgrades), which is the water authority's function, though the office's monitoring may inform the maintenance. It does not include the emergency response (the evacuation, the rescue, the recovery), which is the emergency management authority's function (treated in Chapter 24). The office that expands its scope to include the operation, the allocation, the maintenance, or the response exceeds its mandate and compromises its analytical independence by participating in functions it should only inform.

27.3 The Method

The method of the interface is the structured exchange of data and the coordinated integration of models between the office and the water authority, governed by the inter-agency framework and executed through the waterway data sharing, the model integration, and the event coordination. The method comprises the waterway data ingestion, the flood model execution, the model integration, the flood model delivery, and the joint validation.

The waterway data ingestion is conducted by the office's data pipeline operator (treated in Chapter 15) through the secure data interface. The water authority provides the waterway data in the agreed formats: the channel geometry as the HEC-RAS geometry files (the cross-sections, the reach lengths, the Manning's n values), the levee data as the Shapefiles (the levee alignments, the crest elevations, the design standards), the hydraulic structure data as the rating curves (the discharge rating curves for the gates, the weirs, and the dams), the reservoir data as the time series (the storage levels, the releases, the inflows), and the water allocation data as the monthly reports (the allocations, the diversions, the returns). The operator retrieves the data through the secure data interface, ingests the data into the office's hydrological and hydraulic model database, and logs the ingestion in the office's data ingestion log. The ingestion is monitored for the completeness and the timeliness: if the data is not received by the agreed time (typically the first business day of each month for the standing data, and the daily morning for the reservoir data during the flood season), the operator notifies the office manager, who contacts the authority's data liaison to resolve the delay.

The flood model execution is conducted by the office's hazard specialist (treated in Chapter 17) using the hydrology plug-in (treated in Chapter 10, section 10.2). The HEC-HMS model is

configured with the rainfall (the observed rainfall from the rain gauge network and the radar rainfall estimates, the forecast rainfall from the meteorological service's numerical weather prediction model), the land cover (from the office's land cover data, treated in Chapter 9, section 9.5), the soil moisture (from the NASA SMAP satellite or the in-situ soil moisture sensors), and the baseflow (from the water authority's water allocation data, which provides the baseflow's low-flow conditions). The HEC-HMS run produces the river flows at the catchment outlet and at the intermediate gauge locations. The HEC-RAS 2D model is configured with the LiDAR DEM (the 2-metre, 0.12-metre RMSE DEM treated in Chapter 19, section 19.9), the channel geometry (from the water authority's HEC-RAS geometry files), the levee data (from the water authority's Shapefiles, with the levees represented as the hydraulic barriers in the HEC-RAS model), the hydraulic structure data (from the water authority's rating curves, with the structures represented as the inline structures in the HEC-RAS model), the upstream boundary condition (the HEC-HMS hydrograph), and the downstream boundary condition (the tidal stage or the rating curve, depending on the location). The HEC-RAS run produces the inundation extent and the depth at the 2-metre grid resolution, with the forecast horizon of 48 to 72 hours.

The model integration is conducted through the data interfaces between the office's HEC-RAS model and the water authority's operational models. The integration is bidirectional: the office's HEC-RAS model consumes the water authority's waterway data (the channel geometry, the levee data, the hydraulic structure data, the reservoir data), and the water authority's operational models consume the office's flood predictions (the forecast river stages, the forecast inundation extents). The integration is conducted through the standardized data formats: the HEC-DSS format for the time series data (the hydrographs, the stage hydrographs, the reservoir releases), the HEC-RAS geometry format for the geometric data (the cross-sections, the levees, the structures), and the GeoTIFF format for the raster data (the inundation extents, the depth grids). The integration is tested before each flood season through the joint exercises, with the office and the authority jointly running the models on the historical flood events and verifying the models' consistency.

The flood model delivery is conducted through the secure data interface and the direct communication. The standing flood models (the 100-year, 500-year flood hazard maps, treated in Chapter 19) are delivered annually, with the models delivered to the authority through the secure data interface. The event-driven flood models (the forecast inundation extents during a flood event) are delivered on the event-driven cadence (every 6 to 12 hours during a major event), with the models delivered through the secure data interface for the routine updates and the direct communication for the urgent updates. The direct communication is conducted by the office manager with the authority's operations director, by the phone or the secure messaging, and the communication is logged in the office's event log. The flood model delivery is the operational expression of the office's flood modelling function, and the delivery that meets the authority's needs in the authority's formats and the authority's timelines is the delivery that makes the office's flood models useful for the authority's waterway operations.

The joint validation is conducted after each flood event, with the office and the authority jointly validating the flood models' predictions against the observed data. The validation compares the model's predicted river stages with the observed stages at the gauge locations (the authority's stream gauge network, treated in Chapter 19, section 19.3), the model's predicted inundation extents with the

observed extents (the satellite imagery's observed extents, the field reports' observed extents), and the model's predicted flood durations with the observed durations. The validation results are documented in the model validation report, with the report identifying the model's strengths and weaknesses and the recommendations for the model's improvement. The joint validation is the mechanism that allows the office's flood models to improve over time, and the validation is the condition of the authority's trust in the models.

27.4 The Products

The interface's products are the office's flood risk intelligence products, delivered to the water authority on the standing, event-driven, and monitoring cadence. The products comprise five types, each serving a different phase of the water authority's waterway management cycle.

The flood hazard map is the standing product, produced on the annual or biennial cadence. The map shows the flood hazard zones for the design return periods (the 10-year, 50-year, 100-year, and 500-year flood hazard zones, derived from the hydrological and hydraulic modelling treated in Chapter 19, section 19.3). The map is used by the water authority for the waterway planning (the levee design, the channel capacity assessment, the floodplain management), the reservoir planning (the flood control rules, the storage allocation), and the public education (the publication of the hazard zones for the public awareness). The map carries the Public classification for the general version (intended for the public awareness) and the Internal classification for the operational version (which includes the additional detail on the depths, the velocities, and the arrival times that are restricted to the authority and the operational partners).

The flood inundation map is the event-driven product, produced during a flood event. The map shows the current or forecast inundation extent, updated as the event evolves, with the 6-to-12-hour update cycle during a major event. The current inundation map is derived from the satellite imagery (the Sentinel-1 radar imagery, which can penetrate the cloud cover) and the hydraulic model's simulation of the current conditions. The forecast inundation map is derived from the hydraulic model's forecast run, producing the forecast inundation extent for the next 48 to 72 hours. The map is used by the water authority for the waterway operations (the gate operations, the pump operations, the reservoir releases), the levee monitoring (the identification of the levee sections under the greatest stress), and the emergency coordination (the coordination with the emergency management authority, treated in Chapter 24). The map carries the Internal classification by default, and the Security Confidential classification if the inundation extent affects the critical infrastructure.

The levee and infrastructure monitoring report is the monitoring product, produced on the monthly cadence with the event-driven monitoring during the flood events. The report comprises the infrastructure risk map (showing the monitored levees, floodgates, pumping stations, and channels, with the vulnerability scores colour-coded by priority), the encroachment alerts (identifying the specific encroachments detected during the monthly monitoring cycle), and the maintenance advisories (identifying the infrastructure at risk from the ground conditions and recommending the maintenance actions). The report is used by the water authority for the levee maintenance (the inspections, the repairs, the upgrades), the infrastructure management (the gate and pump maintenance, the channel dredging), and the risk assessment (the levee vulnerability assessment, the

breach potential assessment). The report carries the Internal classification by default, and the Security Confidential classification if the report identifies the critical infrastructure's vulnerabilities.

The model validation report is the analytical product, produced after each flood event. The report compiles the validation results (the comparison of the model's predictions with the observed data), the model's strengths and weaknesses (the aspects of the model that performed well, the aspects that performed poorly), the recommendations for the model's improvement (the parameter adjustments, the geometry updates, the methodological enhancements), and the lessons learned. The report is used by the water authority for the model's credibility assessment (the authority's confidence in the model's predictions for the future events) and the office's model improvement (the model's recalibration, the geometry's update). The report carries the Internal classification, with the Public summary version for the public communication.

The joint flood risk assessment is the integrated product, produced jointly by the office and the water authority on the project-based cadence. The assessment integrates the office's flood hazard maps with the water authority's waterway management plans, producing the integrated assessment of the jurisdiction's flood risk. The assessment comprises the flood hazard (from the office's models), the waterway capacity (from the authority's waterway data), the levee vulnerability (from the office's monitoring and the authority's inspections), the reservoir flood control (from the authority's reservoir data and the office's models), and the integrated risk assessment (the synthesis of all the components into the overall flood risk picture). The assessment is used by the authority for the strategic waterway planning (the long-term levee upgrades, the reservoir expansion, the floodplain restoration) and by the office for the strategic flood model development. The assessment is published jointly under both offices' authorship, with the Public classification for the general version and the Internal classification for the operational version.

27.5 The Consumers

The interface's primary consumer is the water authority itself. The authority's operations division uses the flood inundation map for the waterway operations (the gate operations, the pump operations, the reservoir releases) during the flood events, with the forecast inundation extents informing the operational decisions. The authority's planning division uses the flood hazard map and the joint flood risk assessment for the waterway planning (the levee design, the channel capacity assessment, the floodplain management) and the reservoir planning (the flood control rules, the storage allocation). The authority's maintenance division uses the levee and infrastructure monitoring report for the levee maintenance (the inspections, the repairs, the upgrades) and the infrastructure management (the gate and pump maintenance, the channel dredging). The authority's executive uses the model validation report for the model's credibility assessment and the strategic decisions on the flood risk management.

The emergency management authority is the secondary consumer of the interface's products, particularly the flood inundation map and the joint flood risk assessment. The emergency management authority uses the flood inundation map for the evacuation orders, the resource deployment, and the public warnings, treated in Chapter 24. The emergency management authority uses the joint flood risk assessment for the preparedness phase (the evacuation planning, the resource positioning, the public education). The products that the office delivers to the water authority are also

delivered to the emergency management authority (with the appropriate coordination, since the two authorities are the office's two principal consumers in the flood domain), and the products' consistency across the two consumers is the condition of the two authorities' coordination.

The sectoral regulator is the secondary consumer, particularly for the levee and infrastructure monitoring report. The regulator that regulates the water sector (the dam safety regulator, the levee safety regulator) uses the monitoring report for the compliance verification (the verification that the water authority is maintaining the levees and the infrastructure in accordance with the safety standards). The regulator's use of the monitoring report is treated in Chapter 26 (the sectoral regulator interface), and the office's products are delivered to the regulator through the sectoral regulator interface, not through the water authority interface. The water authority is aware that the office's monitoring report is also delivered to the regulator, and the awareness is the condition of the water authority's trust in the office's monitoring (the authority knows that the office's monitoring is the independent check on the authority's maintenance, and the independent check is the condition of the regulator's confidence in the authority's waterway management).

The planning ministry is the secondary consumer, particularly for the joint flood risk assessment. The planning ministry uses the joint flood risk assessment for the land-use planning (the development restrictions in the flood hazard zones, the floodplain restoration), the infrastructure planning (the road and bridge design in the flood-prone areas, the utility siting), and the environmental planning (the wetland conservation, the riparian restoration). The joint flood risk assessment's integrated picture of the flood risk is the planning ministry's primary input to the flood-related planning decisions, and the assessment's joint authorship (by the office and the water authority) is the condition of the planning ministry's trust in the assessment.

The insurance industry and the public are the tertiary consumers of the interface's products, particularly the flood hazard map. The insurance industry uses the flood hazard map for the flood insurance pricing (the premium calculation, the risk assessment), with the map's Public version providing the general flood hazard zones that the insurance industry uses for the portfolio-level risk assessment. The public uses the flood hazard map for the awareness (the property-level flood risk assessment, the purchase decisions, the insurance decisions). The Public version of the flood hazard map is published on the office's website and the water authority's website, with the open data licence that allows the free use and the redistribution.

The research community and the academic institutions are the tertiary consumers of the interface's products, particularly the model validation report and the joint flood risk assessment. The research community uses the model validation report for the hydrological and hydraulic model research (the model accuracy assessment, the model comparison, the methodological development) and the joint flood risk assessment for the flood risk research (the integrated risk assessment, the climate change impact assessment, the adaptation planning). The products' Public versions are published for the free access, and the Internal versions are shared with the research community through the data-sharing arrangements that the office and the water authority maintain with the academic institutions.

27.6 The Cadence

The interface operates on four cadences: the standing cadence, the event-driven cadence, the monitoring cadence, and the review cadence. The cadences are differentiated by the activity's frequency, the activity's trigger, and the activity's character.

The standing cadence comprises the annual flood hazard map update (the office produces the updated flood hazard map, incorporating the new waterway data from the water authority and the new modelling), the biennial joint flood risk assessment (the office and the water authority jointly produce the integrated assessment), the annual framework review (the office and the water authority jointly review the framework's effectiveness, identifying the issues encountered and the corrective actions), and the periodic coordination meetings (typically quarterly, with the office manager and the authority's planning director reviewing the recent products, the recent operations, the upcoming events, and the framework's operation). The standing cadence is maintained by the office manager, who is the office's point of contact for the authority, and the cadence is the mechanism that maintains the interface's readiness between the events.

The event-driven cadence is activated when a flood event occurs. The activation is initiated by the meteorological service's flood warning (the office's data pipeline operator monitors the meteorological feeds, treated in Chapter 24, section 24.9) or by the water authority's notification (the authority's operations director contacts the office manager to request the event-driven modelling). The activation triggers the office's event-driven response (treated in Chapter 13, section 13.6), which comprises the immediate production of the event-specific flood models (the HEC-HMS and HEC-RAS runs, with the event-specific rainfall and the event-specific waterway data, produced within 2 to 4 hours of the activation), the accelerated waterway data sharing (the authority provides the current reservoir levels, the gate operations, and the levee conditions at the accelerated pace, typically hourly during the major events), the accelerated model integration (the two offices' models exchange the data in the near-real-time, typically every 1 to 3 hours during the major events), and the ongoing production (the regular updates of the event-specific flood models, every 6 to 12 hours during the event's evolution). The event-driven cadence is maintained by the hazard specialist, who produces the event-specific models, and the office manager, who coordinates with the authority's operations director.

The event-driven cadence's intensity scales with the event's severity, with the same scaling treated in Chapter 24, section 24.6 for the emergency management authority interface. A minor flood event (the river stage exceeding the minor flood stage but not the moderate flood stage) may require only the standing models' confirmation and the situational awareness sharing. A moderate flood event (the river stage exceeding the moderate flood stage but not the major flood stage) requires the event-specific models on the 12-hour update cycle. A major flood event (the river stage exceeding the major flood stage) requires the event-specific models on the 6-hour update cycle, the accelerated waterway data sharing on the hourly cycle, and the office manager's direct coordination with the authority's operations director. The intensity's scaling is specified in the event coordination protocol, and the office manager determines the intensity in consultation with the authority's operations director.

The monitoring cadence comprises the monthly levee and infrastructure monitoring (the office's hazard specialist conducts the satellite-based monitoring of the levees, the floodgates, the pumping stations, and the channels, with the monitoring report delivered to the authority by the fifth business

day of each month), the quarterly levee validation (the office and the authority jointly validate the levee data, with the office's monitoring and the authority's in-situ inspections compared and the discrepancies resolved), and the event-driven post-event monitoring (the rapid monitoring of the levee conditions after a flood event, with the breach detection, the damage assessment, and the repair recommendations). The monitoring cadence is maintained by the hazard specialist, who conducts the monitoring, and the production operator, who produces and delivers the monitoring report.

The review cadence comprises the annual framework review (the office and the authority jointly review the framework's effectiveness, identifying the issues encountered and the corrective actions), the post-event review (the joint review of the flood event response, treated in section 27.2 above), and the model validation review (the joint review of the model validation report, with the office's hazard specialist and the authority's hydrologist jointly assessing the model's performance and the improvement recommendations). The review cadence is maintained by the office manager, who schedules and chairs the reviews, and the authority's planning director or operations director, who participates in the reviews.

The cadence is the mechanism that maintains the interface's rhythm and the interface's improvement over time. The standing cadence maintains the regular intelligence provision. The event-driven cadence delivers the intelligence during the events. The monitoring cadence provides the infrastructure monitoring. The review cadence improves the interface after the events. The four cadences together constitute the interface's operational rhythm, and the rhythm is what allows the office to maintain the interface with the water authority across the multiple activities, despite the different frequencies and the different triggers of the activities.

27.7 The Boundary Disputes

The boundary disputes are the disagreements between the office and the water authority over the boundary between the flood risk modelling and the waterway management, and the disputes are the operational reality of the interface. The disputes arise because the boundary is not always clear: the office's flood models may produce recommendations that the water authority considers operational (the office's model may recommend the reservoir release that the authority considers the operational decision, or the office's model may identify the levee vulnerability that the authority considers the operational concern), and the water authority's waterway management may produce data that the office considers analytical (the authority's hydraulic model results may be the analytical input that the office's flood model should consume, or the authority's levee inspection results may be the analytical input that the office's vulnerability assessment should consume). The disputes are resolved through the inter-agency framework's dispute resolution mechanism, which escalates the dispute to the two offices' parent authorities for the resolution.

The boundary disputes are not a sign of the boundary's failure; the disputes are a sign of the boundary's operation, because the boundary that is never disputed is a boundary that is never tested, and the boundary that is never tested may not be a boundary at all. The disputes are the mechanism by which the boundary is tested and clarified, and the disputes are the mechanism by which the boundary is maintained as a functional boundary rather than a territorial one. The office that resolves its boundary disputes through the framework's dispute resolution mechanism maintains the boundary's clarity; an office that resolves its boundary disputes through the informal accommodation or the

territorial assertion erodes the boundary, and the erosion is the failure that the dispute resolution mechanism is designed to prevent.

The common boundary disputes include the reservoir release dispute (the office's flood model recommends the pre-release from the reservoir to create the flood control capacity, and the water authority considers the release decision the operational authority's function), the levee vulnerability dispute (the office's monitoring identifies the levee vulnerability that the authority's inspection has not identified, and the authority disputes the office's identification), the model assumption dispute (the office's flood model uses the assumption that the authority considers inappropriate, such as the levee failure assumption or the gate operation assumption), the data ownership dispute (the authority's waterway data is the authority's data, and the office's use of the data in the public products may be the disputed use), and the public communication dispute (the office's flood hazard map may identify the levee's vulnerability, and the authority may consider the identification the operational concern that should not be publicized).

The dispute resolution mechanism comprises the operational resolution (the office manager and the authority's operations director or planning director attempt to resolve the dispute at the operational level, through the discussion and the negotiation), the framework resolution (if the operational resolution fails, the dispute is escalated to the framework's review committee, with the committee comprising the office's parent ministry and the authority's parent ministry, and the committee resolves the dispute through the formal review), and the political resolution (if the framework resolution fails, the dispute is escalated to the political level, with the minister or the cabinet resolving the dispute through the political decision). The dispute resolution is documented in the office's dispute log, with the dispute's description, the resolution's rationale, and the corrective actions for the future. The dispute resolution is the mechanism that maintains the boundary's clarity over time, and the resolution is the condition of the boundary's continued operation.

The boundary disputes' resolution is guided by the principle that the boundary exists to protect the integrity of both offices' functions. The office's analytical independence is protected by the boundary, and the authority's operational authority is protected by the boundary. The dispute that resolves the boundary in favour of the office's analytical independence at the expense of the authority's operational authority, or in favour of the authority's operational authority at the expense of the office's analytical independence, erodes the boundary. The dispute that resolves the boundary in a way that preserves both offices' integrity maintains the boundary, and the maintenance is the dispute resolution's objective. The disputes are the interface's operational reality, and the disputes' resolution is the interface's continued operation.

27.8 The Comparative Practice

The comparative practice of the water authority interface falls into three principal models: the integrated water management model, the separated-authority model, and the hybrid model. The three models produce different interface geometries, and the geospatial office operates under whichever model the jurisdiction has adopted.

The integrated water management model combines the water authority, the flood risk assessment, and the flood response in a single institution. The model is exemplified by the Dutch Rijkswaterstaat, established in 1798 as the national water authority and responsible for the

Netherlands' water infrastructure (the levees, the dikes, the dams, the storm surge barriers, the pumping stations), the flood risk assessment (the hydraulic modelling, the flood hazard mapping, the levee inspections), and the flood response (the operational management of the water infrastructure during the flood events). The Rijkswaterstaat operates under the Ministry of Infrastructure and Water Management, and it coordinates with the Dutch national water boards (the waterschappen, the sub-national water authorities that manage the regional water infrastructure) and the Deltares research institute (the independent research institute that provides the hydraulic modelling and the flood risk assessment expertise). The geospatial office's interface with the Rijkswaterstaat is the interface between the office and the integrated water authority, with the office providing the satellite-based flood monitoring and the flood hazard mapping that complements the Rijkswaterstaat's hydraulic modelling. The integrated model simplifies the interface, because the water management, the flood risk assessment, and the flood response are in the same institution, and the coordination is internal. The model's risk is that the office's analytical independence may be compromised by the Rijkswaterstaat's operational priorities, and the model's safeguard is the Deltares's independent role as the analytical check on the Rijkswaterstaat's operations.

The separated-authority model maintains the water authority, the flood risk assessment, and the flood response in separate institutions. The model is exemplified by the United States, where the United States Geological Survey (USGS) conducts the flood risk assessment (the stream gauge network, the flood frequency analysis, the flood hazard mapping), the United States Army Corps of Engineers (USACE) manages the water infrastructure (the dams, the levees, the floodgates, with the USACE's own hydraulic modelling and the levee inspections), and the Federal Emergency Management Agency (FEMA) operates the flood response (the flood insurance, the floodplain management, the disaster response). The USGS's National Water Information System provides the stream gauge data to the USACE and the FEMA, the USACE's hydraulic models inform the FEMA's flood insurance rate maps, and the FEMA's floodplain management informs the USACE's levee design. The geospatial office's interface with the US water authorities is the interface between the office and the multiple separated authorities, with the office providing the satellite-based flood monitoring that complements the USGS's stream gauge network, the USACE's hydraulic modelling, and the FEMA's flood hazard mapping. The separated model preserves the offices' institutional independence, but the model requires the more formal coordination mechanisms (the interagency agreements, the data-sharing protocols, the joint exercises) to compensate for the absence of the unified authority. The model's risk is the coordination's complexity, with the multiple offices' coordination requiring the significant interagency effort, and the model's safeguard is the federal framework's established interagency mechanisms (the National Incident Management System, treated in Chapter 24, section 24.8).

The hybrid model combines the elements of the integrated and the separated models, with the water authority, the flood risk assessment, and the flood response partially integrated and partially separated. The model is exemplified by the Australian arrangement, where the Bureau of Meteorology (the BoM, the national weather and climate agency) provides the flood forecasting (the hydrological modelling, the flood warnings), the state water authorities (such as the Melbourne Water, the Sydney Catchment Authority, and the SunWater) manage the water infrastructure (the dams, the levees, the irrigation), and the state emergency services (such as the Victoria State

Emergency Service and the New South Wales State Emergency Service) operate the flood response. The Geoscience Australia (the national geoscience agency, treated in Chapter 28, section 28.8) provides the satellite-based flood monitoring that complements the BoM's flood forecasting and the state water authorities' hydraulic modelling. The hybrid model's interface is the interface between the office and the multiple partially-integrated authorities, with the office's flood models delivered to the BoM, the state water authorities, and the state emergency services. The hybrid model balances the integrated model's simplicity and the separated model's independence, with the national-level integration (the BoM's flood forecasting) and the state-level separation (the state water authorities and the state emergency services).

The Japanese model is a fourth variation, with the Ministry of Land, Infrastructure, Transport and Tourism (MLIT) managing the national water infrastructure (the major rivers, the dams, the levees), the prefectural governments managing the regional water infrastructure (the secondary rivers, the urban drainage), and the Geospatial Information Authority of Japan (GSI, treated in Chapter 24, section 24.8) providing the satellite-based flood monitoring and the elevation data. The Japan Meteorological Agency (JMA) provides the flood forecasting and the flood warnings. The Japanese model is the separated model with the strong ministerial coordination, with the MLIT, the prefectural governments, the JMA, and the GSI coordinating through the Cabinet Office's Disaster Management Bureau during the major flood events. The 2019 Typhoon Hagibis, which caused the flooding of approximately 100 rivers and the breach of approximately 140 levees in the eastern Japan, demonstrated the model's strength and its challenges: the model's strength was the rapid mobilization of the satellite-based flood monitoring by the GSI and the JMA's flood forecasting, and the model's challenge was the coordination across the multiple authorities during the unprecedented event.

The four models share the flood-water boundary: the office produces the flood models, and the water authority manages the waterways. The models differ in the institutional location of the coordination (the internal coordination in the integrated model, the interagency coordination in the separated model, the national-state coordination in the hybrid model, the ministerial coordination in the Japanese model), but the functional boundary is the same in all models. The geospatial office operates under whichever model the jurisdiction has adopted, and the office's interface with the water authority is shaped by the model. The model determines the institutional complexity, the coordination mechanisms, and the boundary disputes' character, but the flood-water boundary, the data-sharing arrangement, and the model integration are the same across the models, and these are the interface's substantive content.

27.9 Worked Example: The Riverbourne Basin Flood Event Coordination with the Water Authority

The flood event coordination is conducted for the Riverbourne Basin flood event of March 2024, treated in Chapter 24, section 24.9. The event is triggered by the intense rainfall on 22-24 March 2024, with the 72-hour rainfall totalling 215 millimetres across the basin (exceeding the 100-year, 72-hour rainfall depth of 185 millimetres, treated in Chapter 19, section 19.9). The Riverbourne at Portsmith reaches the peak stage of 4.2 metres above the flood stage on 25 March 2024 at 04:00. The event inundates approximately 18 square kilometres of the Portsmith metropolitan area, affecting an estimated 38,000 residents and damaging an estimated 4,200 buildings. The coordination with the

water authority is the office's interface with the authority that manages the Riverbourne's waterway infrastructure during the event.

The water authority for the Riverbourne Basin is the Riverbourne Water Authority (RWA), the statutory authority responsible for the Riverbourne's waterway management, including the Riverbourne's main channel (the 220-kilometre channel from the headwaters to the estuary), the three major reservoirs in the upper basin (the Highmoor Reservoir, the Midbourne Reservoir, and the Westbourne Reservoir, with the combined storage capacity of 850 million cubic metres), the levee system along the lower Riverbourne (the 65-kilometre levee system protecting the Portsmith metropolitan area, with the design standard of the 100-year flood), and the floodgates and the pumping stations at the Riverbourne's estuary (the tidal floodgates that prevent the storm surge from propagating up the Riverbourne, and the pumping stations that pump the floodwaters out of the low-lying areas behind the levees).

The coordination begins at 09:30 on 23 March 2024, when the office manager contacts the RWA's operations director to inform the authority of the meteorological service's severe weather warning (the same warning that triggered the office's activation, treated in Chapter 24, section 24.9). The office manager offers the office's event-driven flood modelling, and the operations director accepts the offer and requests the event-specific inundation forecasts for the lower Riverbourne, with the focus on the levee system's vulnerability. The office manager activates the office's event-driven response at 09:45 (the same activation treated in Chapter 24, section 24.9), and the hazard specialist begins the event-specific HEC-HMS and HEC-RAS runs at 10:00.

The waterway data sharing is accelerated at 10:30. The RWA's data liaison provides the current waterway data through the secure data interface: the Highmoor Reservoir's storage level at 92 per cent of capacity (the high storage level that limits the reservoir's flood control capacity), the Midbourne Reservoir's storage level at 88 per cent, the Westbourne Reservoir's storage level at 85 per cent, the current release schedules (the Highmoor releasing 150 cubic metres per second, the Midbourne releasing 120 cubic metres per second, the Westbourne releasing 100 cubic metres per second, with the combined release of 370 cubic metres per second), the levee system's current conditions (the levee inspections conducted on 22 March 2024, with the conditions reported as the satisfactory, with the noted concern at the Levee Section 23 where the levee's crest elevation is at the 100-year design level but the recent subsidence monitoring has shown the 0.15-metre settlement), and the floodgate and the pumping station status (the floodgates closed for the tidal cycle, the pumping stations operating at the 50 per cent capacity). The office's data pipeline operator ingests the data at 10:45, and the data is available for the hazard specialist's HEC-RAS run.

The hazard specialist configures the HEC-RAS model at 11:00 with the event-specific data. The model is configured with the LiDAR DEM (the same 2-metre, 0.12-metre RMSE DEM used in the 100-year flood hazard map, treated in Chapter 19, section 19.9), the channel geometry (the RWA's latest cross-sections, with the 22 March 2024 survey data for the lower Riverbourne), the levee data (the RWA's levee Shapefiles, with the Levee Section 23's crest elevation adjusted to reflect the 0.15-metre settlement), the hydraulic structure data (the RWA's rating curves for the floodgates and the pumping stations, with the current operating status), the upstream boundary condition (the HEC-HMS hydrograph, with the forecast rainfall and the current reservoir releases), and the downstream boundary condition (the forecast storm surge level from the meteorological service). The specialist

runs the model with the levee failure scenarios (the Levee Section 23's failure scenario, with the levee breaching at the design elevation minus the 0.15-metre settlement), and the model produces the inundation extents for the scenarios. The model run takes approximately 90 minutes and produces the forecast inundation extent for 24 March at 06:00 (the forecast peak) and the Levee Section 23's breach inundation extent.

The hazard specialist produces the event-specific flood inundation map at 12:30, with two scenarios: the no-breach scenario (the levee holds, and the inundation is limited to the areas outside the levee system, approximately 12 square kilometres, the same forecast treated in Chapter 24, section 24.9) and the Levee Section 23's breach scenario (the levee breaches, and the inundation extends to the areas behind the levee, approximately 18 square kilometres, with the additional 6 square kilometres inundated by the breach). The specialist's assessment identifies the Levee Section 23's breach scenario as the high-risk scenario, with the breach's probability assessed at 25 per cent (based on the levee's settlement, the forecast flood stage, and the levee's design standard). The specialist's recommendation is that the RWA should immediately inspect the Levee Section 23 and consider the emergency reinforcement.

The office manager delivers the event-specific flood inundation map and the specialist's assessment to the RWA's operations director at 13:00, by the phone and the secure data interface. The operations director reviews the map and the assessment, and the director's response at 13:30 is to dispatch the RWA's levee inspection team to the Levee Section 23 for the immediate inspection. The office manager also delivers the map and the assessment to the emergency management authority's emergency operations centre at 13:30 (the same delivery treated in Chapter 24, section 24.9, with the inundation map informing the evacuation orders issued at 14:30).

The RWA's levee inspection team reports back at 15:00: the Levee Section 23's inspection confirms the 0.15-metre settlement, and the team identifies the early signs of the seepage at the levee's toe (the seepage is the precursor to the levee's failure, with the seepage's progression indicating the levee's internal erosion). The RWA's operations director, on the basis of the inspection, decides at 15:30 to dispatch the emergency levee reinforcement team to the Levee Section 23, with the reinforcement comprising the sandbagging, the geotextile lining, and the temporary crest elevation increase. The reinforcement operations begin at 17:00 and continue through the night, with the RWA's teams working in 12-hour shifts to maintain the reinforcement.

The model integration is conducted at the accelerated pace through the event. The office's hazard specialist updates the HEC-RAS model every 6 hours (at 06:00, 12:00, 18:00, 00:00), with the updated rainfall forecasts, the updated reservoir releases, and the updated levee conditions (the Levee Section 23's reinforcement is incorporated into the model as the increased crest elevation). The RWA's operational models (the authority's own HEC-RAS model of the levee system, and the authority's reservoir operation model) consume the office's flood predictions, and the authority's operational decisions (the reservoir release adjustments, the floodgate operations, the pumping station operations) are informed by the office's predictions. The two offices' models are consistent throughout the event, with the predictions within the 0.20-metre tolerance at the gauge locations.

The event peaks on 25 March at 04:00, with the Riverbourne at Portsmith reaching the peak stage of 4.2 metres above the flood stage. The observed inundation extent is approximately 18 square kilometres, exceeding the no-breach forecast (12 square kilometres) but matching the breach forecast

(18 square kilometres). The Levee Section 23 does not breach, due to the emergency reinforcement, but the observed inundation extent (18 square kilometres) is the same as the breach scenario's extent because the inundation extends through the areas that the no-breach scenario did not forecast (the storm surge's contribution was larger than the forecast, and the rainfall's intensity exceeded the forecast). The office's hazard specialist notes the underestimation and produces the revised inundation map at 06:00 on 25 March (the same revision treated in Chapter 24, section 24.9), incorporating the observed data and the revised forecast.

The RWA's waterway operations through the event include the reservoir release management (the Highmoor Reservoir's release is reduced from 150 to 80 cubic metres per second at 18:00 on 24 March, to limit the downstream peak; the Midbourne and the Westbourne reservoirs' releases are similarly reduced, with the combined release reduced from 370 to 200 cubic metres per second), the floodgate operations (the floodgates are closed for the tidal peaks on 24 and 25 March, with the pumping stations operating at the 100 per cent capacity to pump the floodwaters out of the low-lying areas), and the levee management (the Levee Section 23's reinforcement is maintained through the event, with the reinforcement team monitoring the levee's conditions and the seepage's progression). The RWA's operations are informed by the office's flood predictions throughout the event, with the operations director and the office manager coordinating every 6 hours on the operational decisions.

The event recedes on 26-28 March, and the RWA stands down the emergency operations on 28 March. The office stands down the event-driven response on 28 March (the same stand-down treated in Chapter 24, section 24.9). The post-event review is conducted on 10 April 2024, jointly by the office and the RWA. The review covers the timeline (the activation at 09:30 on 23 March, the waterway data sharing at 10:30, the HEC-RAS run at 11:00, the inundation map at 12:30, the Levee Section 23's inspection at 15:00, the reinforcement at 17:00, the peak at 04:00 on 25 March, the stand-down on 28 March), the products delivered (the event-specific flood inundation map with the two scenarios, the model updates every 6 hours), the issues encountered (the underestimation of the no-breach scenario's inundation extent, the Levee Section 23's settlement that was known but not previously communicated to the office), the lessons learned (the Levee Section 23's settlement should have been communicated to the office in the standing waterway data sharing, and the model's storm surge boundary condition should have been more conservative), and the corrective actions (the RWA will update the waterway data sharing protocol to include the levee settlement data, and the office will update the HEC-RAS model's storm surge boundary condition to use the upper-bound forecast).

The model validation is conducted on 20 April 2024. The validation compares the model's predictions with the observed data: the model's predicted peak stage at the Portsmith gauge is 4.05 metres (the observed peak stage is 4.2 metres, a difference of 0.15 metres, within the 0.30-metre tolerance); the model's predicted inundation extent for the no-breach scenario is 12 square kilometres (the observed extent is 18 square kilometres, an underestimation of 33 per cent, due to the underestimation of the storm surge); the model's predicted inundation extent for the breach scenario is 18 square kilometres (the observed extent is 18 square kilometres, an exact match, but the breach did not occur, and the match is coincidental rather than predictive); the model's predicted arrival time of the peak is 04:30 on 25 March (the observed arrival time is 04:00, a difference of 30 minutes, within the 1-hour tolerance). The validation identifies the model's strengths (the peak stage prediction, the arrival time prediction) and the model's weaknesses (the storm surge contribution, the

breach scenario's interpretation). The validation's recommendations include the improvement of the storm surge boundary condition (the use of the ensemble forecast rather than the deterministic forecast) and the improvement of the breach scenario's communication (the clear statement that the breach scenario is the conditional scenario, not the forecast scenario).

This worked example illustrates the interface's full operation: the activation at 09:30 on 23 March (the office manager's notification of the RWA's operations director), the accelerated waterway data sharing at 10:30 (the reservoir levels, the release schedules, the levee conditions, the floodgate and the pumping station status), the HEC-RAS run at 11:00 (the event-specific configuration, the levee failure scenarios, the 90-minute run), the event-specific inundation map at 12:30 (the no-breach and the breach scenarios, the specialist's assessment of the 25 per cent breach probability), the delivery to the RWA at 13:00, the RWA's levee inspection at 15:00 (the confirmation of the settlement, the identification of the seepage), the RWA's reinforcement decision at 15:30 (the emergency reinforcement dispatched), the reinforcement operations from 17:00 (the 12-hour shifts, the through-the-night operations), the model integration every 6 hours (the consistent predictions within the 0.20-metre tolerance), the peak at 04:00 on 25 March (the 4.2-metre peak stage, the 18-square-kilometre inundation), the revision at 06:00 on 25 March, the stand-down on 28 March, the post-event review on 10 April (the issues identified, the corrective actions), and the model validation on 20 April (the strengths and the weaknesses, the recommendations). The interface's flood-water boundary was maintained throughout: the office produced the flood models, the RWA managed the waterways, and the two functions were coordinated through the secure data interface, the model integration, and the office manager's direct coordination with the operations director. The interface's event coordination was effective: the office's flood models informed the RWA's operational decisions (the levee inspection, the reinforcement, the reservoir releases, the floodgate operations), and the RWA's waterway data informed the office's flood models. The interface's joint validation identified the model's strengths and the weaknesses, and the validation's recommendations will improve the model's performance for the next event. The interface is the foundation of the office's flood risk intelligence delivery to the water authority, and the interface's integrity is the condition of the office's contribution to the jurisdiction's water management and the public's safety during the flood events. The Levee Section 23's emergency reinforcement, informed by the office's flood model and the specialist's assessment, prevented the levee's breach and the catastrophic inundation of the protected area, and the prevention is the interface's contribution to the event's outcome. The interface's value is the connection between the office's flood models and the RWA's waterway operations, and the connection is the reason the interface exists.

Chapter 28 — The National Statistics Agency Interface

28.1 The Interface Defined

The national statistics agency interface is the boundary between the office that produces spatial intelligence and the agency that produces the authoritative statistical record. The geospatial office produces the flood hazard maps, the crop condition assessments, the energy infrastructure monitoring, and the mineral prospectivity studies treated in Chapters 19 through 22, and the national statistics agency produces the decennial census, the ongoing surveys, and the official economic and social statistics. The two offices exchange data under the inter-agency data-sharing framework: the office's spatial products inform the agency's surveys and publications, and the agency's census and survey data inform the office's models and exposure assessments. The two functions are operationally interdependent because each office's products are more accurate and more useful when informed by the other's data, and they are institutionally distinct because the statistical independence to produce the official record and the analytical independence to produce the intelligence are different integrity requirements that should not be combined in a single office.

The interface is governed by the data-sharing arrangement and the independence protection. The data-sharing arrangement is the bilateral framework that specifies the data each office provides, the formats, the delivery channels, the access controls, and the prohibited uses. The arrangement is between two independent institutions, not between a superior and a subordinate: the statistics agency is independent of the executive, and its independence is the condition of its statistical credibility; the geospatial office is part of the executive, and its intelligence is produced for the executive's use. The independence protection is the mechanism that ensures the data-sharing arrangement does not compromise the statistics agency's independence, by prohibiting the use of the agency's data for the executive's policy advocacy and by ensuring the agency's data is shared with the office only for the office's modeling, not for the office's policy recommendations. The independence protection is the structural condition that allows the two offices to cooperate without compromising the statistics agency's credibility, and the protection is treated in detail in section 28.7 below.

The interface's institutional location varies across jurisdictions. In some jurisdictions, the statistics agency is established as an independent statutory authority (with the independence guaranteed by the legislation, the fixed-term leadership, and the independent governance board); in others, the agency is a department within the executive (with the independence protected by the professional norms and the public scrutiny rather than by the statutory guarantees). The geospatial office may sit within the same ministry as the statistics agency (rare), in a different ministry, or in a separate agency. The institutional geometry does not change the functional arrangement: the two offices exchange data under the data-sharing arrangement, and the statistics agency's independence is protected, regardless of the institutional geometry. The institutional geometry does, however, affect the interface's complexity, with the same-ministry variant simplifying the coordination but introducing a risk that the independence is compromised by the ministry's priorities, and the separate-agency variant preserving the independence but requiring more formal coordination mechanisms.

The interface is bilateral in its primary geometry: the geospatial office and the statistics agency are the two parties. The interface is multilateral in its operational expression, because the statistics agency's data and the office's spatial products are also shared with other agencies and the public. The

statistics agency's census data is shared with the planning ministry, the finance ministry, the electoral commission, and the research community. The office's spatial products are shared with the emergency management authority (Chapter 24), the strategic reserve authority (Chapter 25), the sectoral regulators (Chapter 26), the cadastral authority (Chapter 23), and the public. The bilateral interface between the office and the statistics agency is the channel through which the two offices' data flows, and the multilateral sharing is the channel through which the data, once exchanged between the two offices, flows to the other consumers.

28.2 The Scope

The interface covers six activities: census data sharing, survey data sharing, spatial products sharing, spatial-statistical integration, framework maintenance, and joint research. The activities are differentiated by the data flow's direction (statistics-to-office, office-to-statistics, or bidirectional) and by the activity's cadence (decennial, annual, periodic, or project-based).

Census data sharing is the statistics-to-office data flow that provides the office with the demographic foundation for its exposure assessments. The national statistics agency produces the decennial census, which records the population, the households, the demographic characteristics, and the economic characteristics of every person and every household in the jurisdiction, georeferenced to the census geography (the census enumeration areas, the census blocks, the census tracts, and the larger statistical units). The office consumes the census data for its flood exposure assessments (the population exposed to the flood hazard, treated in Chapter 19), its food security assessments (the population dependent on the agricultural production, treated in Chapter 20), its energy infrastructure assessments (the population served by the energy infrastructure, treated in Chapter 21), and its minerals assessments (the population affected by the mineral development, treated in Chapter 22). The census data sharing is conducted decennially, with the data delivered to the office within 12 months of the census's completion, and annually, with the intercensal estimates delivered to the office as the agency produces them.

Survey data sharing is the statistics-to-office data flow that provides the office with the economic and social foundation for its models. The statistics agency produces the ongoing surveys, including the agricultural production surveys (the cropped area, the yield, the production, by crop and by region), the industrial production surveys (the mineral production, the energy production, the manufacturing production), the labor force surveys (the employment, the unemployment, the wages), the household surveys (the income, the expenditure, the consumption), and the special surveys (the environmental surveys, the infrastructure surveys, the disaster impact surveys). The office consumes the survey data for its crop models (the agricultural production surveys provide the planted area and the historical yield for the yield forecasting, treated in Chapter 20, section 20.3), its energy infrastructure models (the industrial production surveys provide the energy production and consumption data), and its mineral models (the industrial production surveys provide the mineral production data for the resource estimates). The survey data sharing is conducted on the surveys' cadence (monthly for the labor force surveys, quarterly for the agricultural production surveys, annually for the industrial production surveys), with the data delivered to the office within 30 days of the surveys' publication.

Spatial products sharing is the office-to-statistics data flow that provides the agency with the spatial foundation for its surveys and publications. The office produces the land cover data (from the satellite imagery classification, treated in Chapter 9, section 9.5), the infrastructure data (from the infrastructure monitoring, treated in Chapter 21), the administrative boundaries data (from the cadastral authority interface, treated in Chapter 23), and the thematic data (the flood hazard data, the crop condition data, the energy infrastructure data, the mineral prospectivity data, each treated in the respective domain chapters). The statistics agency consumes the spatial products for its survey design (the land cover data for the agricultural survey sampling frames, the administrative boundaries for the census geography), its survey estimation (the infrastructure data for the economic surveys' stratification, the land cover data for the agricultural surveys' area-frame estimation), and its publication (the thematic data for the statistical publications' spatial context, the administrative boundaries for the statistical maps). The spatial products sharing is conducted on the products' cadence (annually for the land cover, monthly for the crop condition, event-driven for the flood hazard), with the products delivered to the agency through the secure data interface.

Spatial-statistical integration is the bidirectional activity that produces the integrated products combining the spatial and the statistical data. The integration is conducted for the products that neither office could produce alone: the population exposure maps (the census data joined to the flood hazard maps, producing the maps of the population exposed to the flood hazard), the agricultural production estimates (the survey data joined to the crop condition data, producing the spatially-explicit yield forecasts), the energy consumption estimates (the survey data joined to the energy infrastructure data, producing the spatially-explicit consumption maps), and the disaster impact estimates (the census data joined to the flood inundation maps, producing the population and building exposure estimates treated in Chapter 24, section 24.4). The integration is conducted jointly by the office and the statistics agency, with each office providing its data and its analytical expertise, and the integrated products are published jointly under both offices' authorship.

Framework maintenance is the standing maintenance of the data-sharing arrangement between the data exchanges. The framework comprises the data-sharing agreement (which specifies the data each office provides, the formats, the delivery channels, the access controls, and the prohibited uses), the confidentiality protocol (which specifies the statistical confidentiality protections applied to the shared data), the integration protocol (which specifies the joint production procedures for the integrated products), and the review cycle (which specifies the annual review of the framework's effectiveness). The framework is maintained by the office manager and the statistics agency's data liaison, and the framework is reviewed annually, after each decennial census, and after each major methodological change in either office's products.

Joint research is the project-based activity that produces the research outcomes combining the spatial and the statistical expertise. The joint research may include the methodological development (the new methods for the spatial-statistical integration, the new methods for the small-area estimation, the new methods for the data fusion), the experimental products (the experimental population grids, the experimental high-resolution poverty maps, the experimental real-time population estimates), and the publications (the joint research papers, the joint conference presentations, the joint methodological reports). The joint research is conducted by the office's technical lead and the statistics agency's methodological director, with the research outcomes

published jointly under both offices' authorship and the methodological findings incorporated into the offices' production procedures.

The scope does not include the statistical production (the census, the surveys, the official statistics), which is the statistics agency's function. It does not include the policy analysis (the policy recommendations, the policy advocacy), which is the executive's function and the prohibited use under the independence protection. It does not include the public communication of the statistics (the statistical publications, the press releases), which is the statistics agency's function. The office that expands its scope to include the statistical production, the policy analysis, or the public communication exceeds its mandate and compromises both its analytical independence and the statistics agency's statistical independence.

28.3 The Method

The method of the interface is the structured exchange of data between the office and the statistics agency, governed by the data-sharing arrangement and executed through the secure data interface, the confidentiality protocols, and the integration procedures. The method comprises the data delivery, the data ingestion, the confidentiality protection, the spatial-statistical integration, and the joint production.

The data delivery is conducted through the secure data interface, which is a bidirectional web service that allows each office to push the data to the other office's platform and to pull the data from the other office's archive. The interface operates over the HTTPS protocol with the mutual TLS authentication, ensuring that the data is transmitted securely and that only the authorized systems may access the interface. The interface is monitored for the availability and the performance, and the office's data pipeline operator is responsible for the interface's technical operation on the office's side, with the statistics agency's data operations team responsible on the agency's side. The data delivery is conducted on the cadence specified in the data-sharing agreement (decennially for the census, annually for the intercensal estimates, monthly for the labor force surveys, quarterly for the agricultural production surveys, annually for the land cover), and the delivery is logged in the office's data ingestion log for the audit trail.

The data ingestion is conducted by the office's data pipeline operator (treated in Chapter 15) for the statistics-to-office data flow, and by the statistics agency's data operations team for the office-to-statistics data flow. The ingestion comprises the data format verification (the data is in the agreed format, typically a CSV, an XML, or a GeoPackage, with the agreed schema), the data quality verification (the data is complete, the data is consistent, the data is within the expected ranges), the data ingestion into the office's database (the data is loaded into the appropriate tables, with the foreign keys to the existing data), and the ingestion logging (the data's source, the data's date, the data's record count, the data's quality flags). The ingestion is monitored for the failures, and the failures are resolved through the inter-agency coordination between the office's data pipeline operator and the statistics agency's data operations team.

The confidentiality protection is applied to the statistics-to-office data flow, because the census and survey data contain personally identifiable information. The protection is governed by the statistical confidentiality framework, which is the framework that the statistics agency operates under and that the office must comply with when it receives the data. The framework comprises the de-

identification (the personally identifiable information is removed or anonymized before the data is shared), the aggregation (the data is aggregated to the level that prevents the re-identification, typically the census enumeration area or the census block), the suppression (the small cells are suppressed to prevent the re-identification, typically the cells with fewer than 10 individuals or 5 households), and the access control (the data is restricted to the authorized personnel, with the access logged for the audit trail). The office's data pipeline operator applies the confidentiality protection on the ingestion, the office's specialists apply the protection on the analysis (the analytical products do not reveal the individual-level data), and the office's production operator applies the protection on the products (the products do not include the individual-level data or the small cells).

The spatial-statistical integration is conducted through the geospatial join, which is the operation that combines the spatial data (the flood hazard map, the crop condition map, the infrastructure map) with the statistical data (the census data, the survey data) based on the common geographic key. The common geographic key is typically the census enumeration area or the census block, which is the smallest census geography that the statistics agency publishes. The office's spatial products are aggregated or disaggregated to the census geography, and the statistics agency's data is joined to the spatial products through the census geography's identifier. The integration produces the spatially-explicit statistical products (the population exposure map, the building exposure map, the agricultural production map) and the statistically-informed spatial products (the flood hazard map with the population exposure, the crop condition map with the production estimates). The integration is conducted by the office's hazard specialist or the resource economics specialist, with the statistics agency's methodological director providing the statistical expertise for the integration's methodological soundness.

The joint production is conducted for the integrated products that are published jointly under both offices' authorship. The joint production comprises the joint planning (the office and the statistics agency jointly plan the integrated product, with the scope, the methodology, the timeline, and the publication), the joint production (each office produces its contribution, with the office producing the spatial component and the statistics agency producing the statistical component), the joint integration (the two components are integrated, with the integration conducted by the office's specialist and the statistics agency's methodological director), the joint review (the integrated product is reviewed by both offices' technical leads and authorized by both offices' office managers or directors), and the joint publication (the integrated product is published under both offices' authorship, with the joint copyright and the joint disclaimer). The joint production is conducted on the project-based cadence, with each project's timeline determined by the project's scope and the joint production's complexity.

28.4 The Products

The interface's products are the data and the integrated products exchanged between the office and the statistics agency, and the integrated products published jointly. The products comprise six types, differentiated by the data flow's direction and the product's character.

The census data is the statistics-to-office product, delivered decennially with the annual intercensal updates. The data comprises the population data (the total population, the age, the sex, the household composition, georeferenced to the census enumeration area), the household data (the

housing tenure, the dwelling type, the household facilities, georeferenced to the census enumeration area), the demographic data (the education, the occupation, the income, georeferenced to the census enumeration area), and the economic data (the employment, the industry, the occupation, georeferenced to the census enumeration area). The census data is the demographic foundation of the office's exposure assessments, and the data is delivered under the statistical confidentiality framework with the de-identification, the aggregation, and the access control.

The survey data is the statistics-to-office product, delivered on the surveys' cadence. The data comprises the agricultural production surveys (the cropped area, the yield, the production, by crop and by region, delivered quarterly), the industrial production surveys (the mineral production, the energy production, the manufacturing production, delivered annually), the labor force surveys (the employment, the unemployment, the wages, delivered monthly), and the household surveys (the income, the expenditure, the consumption, delivered annually). The survey data is the economic and social foundation of the office's models, and the data is delivered under the statistical confidentiality framework with the same protections as the census data.

The land cover data is the office-to-statistics product, delivered annually. The data is produced by the office's data pipeline operator through the satellite imagery classification (treated in Chapter 9, section 9.5), and it comprises the land cover classes (the urban, the agricultural, the forest, the water, the barren, the wetland), the land cover changes (the conversions between the classes, with the change year), and the classification accuracy (the confusion matrix, the overall accuracy, the class-level accuracies). The land cover data is consumed by the statistics agency for the agricultural survey sampling frames (the cropped area frames), the environmental surveys' stratification (the land cover strata), and the statistical publications' spatial context (the land cover maps). The land cover data carries the Public classification, because the data is derived from the public satellite imagery and does not include the security-sensitive or the commercially-sensitive information.

The administrative boundaries data is the office-to-statistics product, delivered annually with the incremental updates. The data is consumed by the statistics agency for the census geography (the census enumeration areas, the census blocks, the census tracts, the larger statistical units), the survey sampling (the primary sampling units, the strata), and the statistical publications (the statistical maps). The administrative boundaries data is obtained by the office from the cadastral authority (treated in Chapter 23), and the office serves as the conduit for the data's delivery to the statistics agency, with the cadastral authority's authorization. The data carries the Internal classification, because the boundaries may include the disputed or the security-sensitive information.

The thematic data is the office-to-statistics product, delivered on the products' cadence. The data comprises the flood hazard data (the flood hazard zones, delivered annually), the crop condition data (the condition assessments, delivered weekly during the growing season), the energy infrastructure data (the infrastructure risk maps, delivered monthly, Security Confidential), and the mineral prospectivity data (the prospectivity maps, delivered project-based, Confidential). The thematic data is consumed by the statistics agency for the statistical publications' spatial context (the flood hazard maps in the disaster statistics publications, the crop condition maps in the agricultural statistics publications), the survey design (the flood hazard data for the disaster impact surveys' stratification), and the joint integrated products (the flood hazard data joined to the census data for the population exposure maps). The thematic data is delivered under the classification protections that apply to each

product, with the Security Confidential and the Confidential products delivered through the restricted channels.

The integrated products are the joint products, published under both offices' authorship. The products comprise the population exposure maps (the census data joined to the flood hazard maps, published jointly as the official population exposure assessments), the agricultural production estimates (the survey data joined to the crop condition data, published jointly as the official spatially-explicit yield estimates), the energy consumption estimates (the survey data joined to the energy infrastructure data, published jointly as the official spatially-explicit consumption estimates), and the disaster impact estimates (the census data joined to the flood inundation maps, published jointly as the official disaster impact assessments). The integrated products carry the classification of the most-restricted input, with the Security Confidential and the Confidential products published under the restricted access and the Public products published for the public access.

28.5 The Consumers

The interface's primary consumers are the geospatial office and the statistics agency themselves. The office consumes the census and survey data for its models and exposure assessments, and the agency consumes the spatial products for its surveys and publications. The two offices are the interface's primary consumers because the interface is the channel through which the two offices' data flows, and the data flow is the mechanism that allows each office to produce better products than either could produce alone.

The emergency management authority is the secondary consumer of the interface's integrated products, particularly the population exposure maps and the disaster impact estimates. The authority uses the population exposure maps (the census data joined to the flood hazard maps) for the evacuation planning, the resource positioning, and the public education, treated in Chapter 24. The authority uses the disaster impact estimates (the census data joined to the flood inundation maps) for the response phase, the evacuation orders, the resource deployment, and the public warnings. The integrated products provide the authority with the spatially-explicit population and building exposure that the authority's common operating picture requires, and the products' statistical credibility (derived from the statistics agency's census data) is the condition of the authority's trust.

The strategic reserve authority is the secondary consumer of the interface's integrated products, particularly the agricultural production estimates. The reserve authority uses the spatially-explicit yield estimates (the survey data joined to the crop condition data) for the procurement advisories, treated in Chapter 25. The integrated products provide the reserve authority with the yield forecasts that are statistically grounded in the agricultural production surveys, and the products' statistical credibility is the condition of the reserve authority's trust in the procurement advisories.

The sectoral regulators are the secondary consumers of the interface's integrated products. The energy regulator uses the energy consumption estimates (the survey data joined to the energy infrastructure data) for the regulatory policy and the compliance verification, treated in Chapter 26. The minerals regulator uses the mineral production estimates (the survey data joined to the mineral prospectivity data) for the regulatory policy and the licence management. The integrated products provide the regulators with the spatially-explicit statistical products that the regulators' regulatory functions require, and the products' statistical credibility is the condition of the regulators' trust.

The planning ministry is the secondary consumer of the interface's integrated products, particularly the population exposure maps and the land cover data. The planning ministry uses the population exposure maps for the land-use planning (the development restrictions in the flood hazard zones, the population allocation across the regions) and the land cover data for the environmental planning (the conservation areas, the urban growth boundaries, the agricultural preservation). The integrated products provide the planning ministry with the spatially-explicit statistical products that the ministry's planning function requires, and the products' statistical credibility is the condition of the ministry's trust.

The research community and the public are the tertiary consumers of the interface's integrated products, particularly the Public products. The research community uses the Public integrated products (the population exposure maps, the agricultural production estimates) for the academic research, the policy analysis, and the public-interest advocacy. The public uses the Public integrated products for the awareness, the education, and the civic engagement. The integrated products' Public versions are published on the statistics agency's website and the office's website, with the open data licences that allow the free use and the redistribution. The Public versions do not include the security-sensitive, the commercially-sensitive, or the personally-identifiable information, and the publication is governed by the statistical confidentiality framework and the office's classification regime.

28.6 The Cadence

The interface operates on four cadences: the decennial cadence, the annual cadence, the ongoing cadence, and the project-based cadence. The cadences are differentiated by the data flow's frequency and the activity's character.

The decennial cadence is the cadence of the census data sharing. The census is conducted every 10 years (in the years ending in 0 in most jurisdictions, with some variations), and the census data is delivered to the office within 12 months of the census's completion. The decennial cadence is the most consequential of the interface's cadences, because the census data is the demographic foundation of the office's exposure assessments, and the data's update every 10 years is the event that requires the office to update its exposure assessments' demographic foundation. The decennial cadence is maintained by the office's data pipeline operator, who ingests the census data, and the office's hazard specialist, who updates the exposure assessments' demographic foundation.

The annual cadence is the cadence of the annual products and the framework review. The annual products comprise the annual intercensal population estimates (delivered by the statistics agency to the office), the annual land cover data (delivered by the office to the statistics agency), the annual administrative boundaries data (delivered by the office to the statistics agency), the annual industrial production surveys (delivered by the statistics agency to the office), and the annual household surveys (delivered by the statistics agency to the office). The annual framework review is conducted by the office manager and the statistics agency's data liaison, with the review covering the framework's operation, the data exchanges' completeness and timeliness, the confidentiality protection's effectiveness, and the corrective actions. The annual cadence is maintained by the office's data pipeline operator (for the data ingestion), the office's hazard specialist and resource economics specialist (for the data consumption), and the office manager (for the framework review).

The ongoing cadence is the cadence of the ongoing data exchanges. The ongoing data exchanges comprise the monthly labor force surveys (delivered by the statistics agency to the office), the quarterly agricultural production surveys (delivered by the statistics agency to the office), the weekly crop condition data (delivered by the office to the statistics agency during the growing season), the monthly energy infrastructure data (delivered by the office to the statistics agency, Security Confidential), and the event-driven flood inundation maps (delivered by the office to the statistics agency during the flood events). The ongoing cadence is maintained by the office's data pipeline operator, who manages the data exchanges on the ongoing cadence, and the specialists, who consume the data on the ongoing cadence.

The project-based cadence is the cadence of the joint research and the integrated products. The joint research projects are conducted on the timelines determined by the projects' scope, typically 6 to 24 months for the methodological development, 3 to 12 months for the experimental products, and 12 to 36 months for the joint publications. The integrated products are produced on the timelines determined by the products' scope, typically 3 to 6 months for the population exposure maps (annual), 6 to 12 months for the agricultural production estimates (annual), and 1 to 3 months for the disaster impact estimates (event-driven). The project-based cadence is maintained by the office's technical lead (for the joint research) and the office's hazard specialist or the resource economics specialist (for the integrated products), with the statistics agency's methodological director providing the statistical expertise.

The cadence is the mechanism that maintains the interface's rhythm and the interface's improvement over time. The decennial cadence provides the demographic foundation. The annual cadence provides the annual products and the framework review. The ongoing cadence provides the ongoing data exchanges. The project-based cadence provides the joint research and the integrated products. The four cadences together constitute the interface's operational rhythm, and the rhythm is what allows the office to maintain the interface with the statistics agency across the multiple data flows and the multiple activities, despite the different frequencies and the different scopes of the activities.

28.7 The Independence Protection

The independence protection is the mechanism by which the data-sharing arrangement protects the statistics agency's independence from the executive, and the mechanism is the interface's distinctive requirement. The statistics agency's independence is the condition of its statistical credibility, because the statistics that are not independent of the executive are not statistics but propaganda, and the propaganda's publication under the statistical authority's name undermines both the statistics agency's credibility and the public's trust in the government's information. The independence protection is the mechanism that ensures the data-sharing arrangement does not compromise the statistics agency's independence, and the protection is maintained through the data-sharing agreement's prohibitions, the confidentiality framework, and the institutional safeguards.

The data-sharing agreement's prohibitions are the first layer of the independence protection. The agreement prohibits the use of the statistics agency's data for the executive's policy advocacy: the office may use the census and survey data for its models and exposure assessments, but the office may not use the data to support the executive's policy recommendations or the executive's political

messaging. The office's products that draw on the statistics agency's data must be analytical (describing the conditions and the trends) and not advocative (recommending the policies or supporting the political positions). The prohibition is the mechanism that prevents the office's use of the statistics agency's data from being misconstrued as the statistics agency's endorsement of the executive's policies, and the prohibition is the condition of the statistics agency's independence.

The confidentiality framework is the second layer of the independence protection. The framework protects the individuals' and the businesses' privacy, and the protection is the condition of the public's trust in the statistics agency's data collection. The public's trust is the condition of the public's cooperation with the census and the surveys, and the public's cooperation is the condition of the data's completeness and accuracy. The office that discloses the statistics agency's data, or that uses the data in a way that allows the re-identification, undermines the public's trust, and the undermining is the failure that the confidentiality framework is designed to prevent. The framework is maintained by the office's data pipeline operator (on the ingestion), the office's specialists (on the analysis), and the office's production operator (on the products), with each office's role specified in the data-sharing agreement.

The institutional safeguards are the third layer of the independence protection. The safeguards comprise the statistics agency's statutory independence (the agency is established by the legislation, with the fixed-term leadership, the independent governance board, and the protected budget), the agency's professional norms (the agency's staff are bound by the professional ethics, the statistical principles, and the confidentiality obligations), and the public scrutiny (the agency's products are published, the agency's methods are documented, and the agency's decisions are subject to the public review). The institutional safeguards are the mechanism that ensures the statistics agency's independence is maintained across the changes of leadership, the changes of budget, and the changes of political priority, and the safeguards are the condition of the agency's long-term credibility.

The independence protection is consistent with the international statistical frameworks. The United Nations Fundamental Principles of Official Statistics, adopted by the United Nations Statistical Commission in 1994 and endorsed by the United Nations General Assembly in 2014, establish the independence of the official statistics as the first principle: "Official statistics provide an indispensable element in the information system of a democratic society, serving the government, the economy and the public with data about a wide range of activities. To retain trust in official statistics, the statistical agencies need to decide according to strictly professional considerations, including scientific principles and professional ethics, on the methods and procedures for the collection, processing, storage and presentation of statistical data." The European Statistics Code of Practice, adopted by the European Statistical System in 2005 and revised in 2011 and 2017, establishes the independence of the European statistical system as the first principle: "Professional independence of statistical authorities from other policy, regulatory or administrative departments and bodies, as well as from private sector operators, ensures the credibility of European Statistics." The data-sharing arrangement's independence protection is aligned with these frameworks, and the office's compliance with the protection is the condition of the office's continued access to the statistics agency's data.

The independence protection is the interface's distinctive requirement, and the requirement is what distinguishes the national statistics agency interface from the other sectoral interfaces treated in this Part. The emergency management authority interface (Chapter 24), the strategic reserve authority

interface (Chapter 25), the sectoral regulator interface (Chapter 26), and the cadastral authority interface (Chapter 23) are all interfaces between the office and the other executive agencies, where the institutional arrangements are governed by the executive's coordination and the inter-agency frameworks. The national statistics agency interface is the interface between the office and an independent institution, where the institutional arrangements are governed by the statistics agency's statutory independence and the independence protection. The office that respects the independence protection draws on the statistics agency's data without compromising the agency's independence; an office that does not respect the protection compromises the agency's independence, and the compromise undermines both the agency's credibility and the office's continued access to the data.

28.8 The Comparative Practice

The comparative practice of the national statistics agency interface varies by jurisdiction, but the models share the data-sharing arrangement and the independence protection. The comparative practice is examined in three jurisdictions: the United States, Australia, and the Netherlands, with the three jurisdictions representing three variants of the institutional arrangement and the data-sharing framework.

The United States model is exemplified by the U.S. Census Bureau and the U.S. Geological Survey (USGS). The Census Bureau, established in 1902 as a permanent agency within the Department of Commerce, produces the decennial census (mandated by Article I, Section 2 of the U.S. Constitution), the American Community Survey (the annual survey that replaced the long-form census in 2005), and the economic surveys (the Economic Census, the Census of Governments, the annual economic surveys). The Bureau's independence is protected by the Title 13 of the U.S. Code, which establishes the Bureau's confidentiality protections and the Bureau's professional independence, and by the Bureau's governance arrangements (the Director's five-year term, the Office of Management and Budget's statistical policy oversight, and the Congressional oversight). The USGS, established in 1879 within the Department of the Interior, produces the satellite imagery (through the Landsat programme, treated in Chapter 8), the land cover data (through the National Land Cover Database), the flood hazard maps (through the National Flood Hazard Layer), and the mineral resource assessments. The data-sharing arrangement between the Census Bureau and the USGS is governed by the Federal Statistical System's data-sharing framework, with the Census Bureau providing the census and survey data to the USGS for the exposure assessments and the USGS providing the land cover and the flood hazard data to the Census Bureau for the survey design and the publications. The arrangement's independence protection is maintained through the Title 13 confidentiality protections, the Federal Statistical System's governance, and the offices' professional norms.

The Australian model is exemplified by the Australian Bureau of Statistics (ABS) and the Geoscience Australia (GA). The ABS, established in 1905 as the Commonwealth Bureau of Census and Statistics and renamed in 1975, produces the Census of Population and Housing (conducted every five years, more frequently than the decennial census in most jurisdictions), the economic surveys (the Agricultural Census, the Economic Census, the monthly Labour Force Survey), and the social surveys (the National Health Survey, the General Social Survey). The ABS's independence is protected by the Australian Bureau of Statistics Act 1975, which establishes the ABS's independence

from the executive and the Australian Statistician's statutory authority, and by the ABS's governance arrangements (the Australian Statistician's independent authority, the ABS's independent advisory council, and the Parliamentary oversight). Geoscience Australia, established in 2001 from the merger of the Australian Geological Survey Organisation and the Australian Surveying and Land Information Group, produces the satellite imagery (through the Digital Earth Australia programme, treated in Chapter 8), the land cover data, the flood hazard maps, and the mineral resource assessments. The data-sharing arrangement between the ABS and the GA is governed by the Australian data-sharing framework, with the ABS providing the census and survey data to the GA for the exposure assessments and the GA providing the land cover and the flood hazard data to the ABS for the survey design and the publications. The arrangement's independence protection is maintained through the ABS Act's confidentiality protections, the Australian data-sharing framework's governance, and the offices' professional norms.

The Netherlands model is exemplified by the Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS) and the Netherlands' Cadastre, Land Registry and Mapping Agency (Kadaster). The CBS, established in 1899, produces the Dutch statistical record, including the population statistics (based on the municipal population registers rather than a separate census, the Netherlands having transitioned to a register-based census in 1991), the economic statistics, and the social statistics. The CBS's independence is protected by the Statistics Netherlands Act 1996 (the Wet op het Centraal bureau en de Centrale commissie voor de statistiek), which establishes the CBS's independence and the Central Statistics Commission's governance, and by the European Statistics Code of Practice (the CBS being part of the European Statistical System). The Kadaster, established in 1838, produces the Dutch cadastre (treated in Chapter 23), the land registry, and the large-scale topographic maps (the Basisregistratie Topografie, BTop10). The data-sharing arrangement between the CBS and the Kadaster is governed by the Dutch Basisregistraties framework (the system of the national base registries, including the Basisregistratie Personen, the Basisregistratie Adressen en Gebouwen, and the Basisregistratie Topografie), with the CBS providing the population and the economic statistics to the Kadaster and the Kadaster providing the cadastral and the topographic data to the CBS. The arrangement's independence protection is maintained through the CBS Act's confidentiality protections, the Basisregistraties framework's governance, and the European Statistics Code of Practice.

The three models share the data-sharing arrangement and the independence protection, but they differ in the institutional arrangements and the data-sharing frameworks. The U.S. model is the federal-statistical-system variant, with the Census Bureau and the USGS in different departments and the data-sharing governed by the federal statistical framework. The Australian model is the national-statistical-agency variant, with the ABS and the GA in different portfolios and the data-sharing governed by the national data-sharing framework. The Netherlands model is the register-based variant, with the CBS and the Kadaster operating under the Basisregistraties framework, and the data-sharing conducted through the national base registries rather than through the bilateral data-sharing arrangements. The Netherlands model is the example of the integrated variant, where the data-sharing is institutionalized in the national base registries and the bilateral arrangements are simplified by the registries' interoperability. The three models also share the spatial-statistical integration challenge, with the integration conducted through the geospatial join on the common geographic key, and the

integration producing the spatially-explicit statistical products that neither office could produce alone.

The comparative practice shows that the national statistics agency interface is the interface between the office and an independent institution, and the interface's distinctive requirement is the independence protection. The geospatial office operates under whichever model the jurisdiction has adopted, and the office's interface with the statistics agency is shaped by the model. The model determines the institutional complexity, the data-sharing framework, and the integration method, but the data-sharing arrangement and the independence protection are the same across the models, and these are the interface's substantive content.

28.9 Worked Example: The Portsmith Flood Exposure Assessment

The flood exposure assessment is conducted for the Portsmith metropolitan area, the coastal city of 250,000 inhabitants at the mouth of the Riverbourne treated in Chapter 19, section 19.9. The city's flood hazard is the 100-year flood hazard map produced in Chapter 19, section 19.9, which integrates the riverine flood hazard (from the Riverbourne) and the coastal flood hazard (from the storm surge), and which shows the 100-year flood hazard zones with the depths and the velocities. The exposure assessment joins the flood hazard map to the census data to produce the population and building exposure estimates for the 100-year flood hazard, and the assessment is the joint product of the geospatial office and the national statistics agency.

The assessment begins on 5 February 2024, when the office manager and the statistics agency's data liaison jointly plan the integrated product. The planning meeting, conducted in the office's conference room, agrees on the scope (the 100-year flood hazard map for the Portsmith metropolitan area, joined to the 2020 census data), the methodology (the geospatial join on the census enumeration area, with the area-weighted aggregation for the enumeration areas that straddle the flood hazard zones), the timeline (the assessment to be produced by 31 March 2024, 8 weeks from the planning meeting), and the publication (the joint publication under both offices' authorship, with the Public classification for the general version and the Internal classification for the operational version).

The office's hazard specialist produces the spatial component on 12-16 February 2024. The specialist retrieves the 100-year flood hazard map (produced in Chapter 19, section 19.9) from the office's archive, verifies the map's coverage (the Portsmith metropolitan area, 18 square kilometres), and prepares the map for the geospatial join: the flood hazard zones are dissolved to the inundation extent (the area inundated by the 100-year flood, with the depths), and the inundation extent is intersected with the census enumeration areas (1,247 enumeration areas in the Portsmith metropolitan area, with the average area of 1.44 hectares). The intersection produces the 1,247 enumeration-area-level inundation statistics: the inundated area, the inundation depth (the area-weighted average), and the inundation percentage (the ratio of the inundated area to the enumeration area's total area). The specialist submits the spatial component to the office's technical lead for the validation on 19 February 2024.

The statistics agency's methodological director produces the statistical component on 19-23 February 2024. The director retrieves the 2020 census data for the Portsmith metropolitan area from the agency's database, with the data comprising the population (the total population, the age, the sex, the household composition), the households (the housing tenure, the dwelling type, the household

facilities), and the buildings (the building type, the building age, the building value) for each of the 1,247 enumeration areas. The director applies the statistical confidentiality framework: the data is de-identified (the personally identifiable information is removed), the data is aggregated (the data is at the enumeration-area level, the smallest census geography that the agency publishes), and the small cells are suppressed (the enumeration areas with fewer than 10 individuals or 5 households are suppressed, with 14 enumeration areas suppressed in the Portsmith metropolitan area). The director submits the statistical component to the agency's quality assurance team for the validation on 26 February 2024.

The joint integration is conducted on 1-5 March 2024 by the office's hazard specialist and the statistics agency's methodological director, working together in the office's advisory production environment. The integration is the geospatial join of the spatial component (the enumeration-area-level inundation statistics) and the statistical component (the enumeration-area-level census data) on the common geographic key (the census enumeration area's identifier). The join produces the 1,247 enumeration-area-level exposure records: the inundated area, the inundation depth, the inundation percentage, the population, the households, and the buildings for each enumeration area. The area-weighted aggregation is applied for the enumeration areas that straddle the flood hazard zones: the population and the buildings in the partially-inundated enumeration areas are estimated as the total population and the buildings multiplied by the inundation percentage, with the assumption that the population and the buildings are uniformly distributed within the enumeration area. The assumption's limitation is documented in the assessment's methodology section, and the limitation's implication for the assessment's accuracy is discussed.

The integration produces the exposure estimates for the 100-year flood hazard in the Portsmith metropolitan area. The total population exposed to the 100-year flood hazard is 38,420 (15.4 per cent of the Portsmith metropolitan area's population of 250,000), with the exposure concentrated in the central business district's low-lying areas (12,800 exposed), the residential neighbourhoods along the Riverbourne's left bank (18,200 exposed), and the industrial estate at the basin's outlet (7,420 exposed). The total households exposed are 14,850 (15.0 per cent of the metropolitan area's households), and the total buildings exposed are 11,200 (13.8 per cent of the metropolitan area's buildings), with the buildings categorized by the type: 8,400 residential, 1,800 commercial, 650 industrial, and 350 public. The critical infrastructure exposed includes 2 hospitals, 4 schools, 1 substation, 1 water-treatment plant, and 2 hazardous-materials sites (the same critical infrastructure identified in the worked example in Chapter 24, section 24.9, on the March 2024 flood event).

The joint review is conducted on 8-12 March 2024. The integrated product is reviewed by the office's technical lead (for the spatial component's accuracy and the integration's methodological soundness) and the statistics agency's quality assurance team (for the statistical component's accuracy and the confidentiality protection's effectiveness). The review identifies two issues: (1) the area-weighted aggregation's assumption of the uniform distribution is inadequate for the enumeration areas with the concentrated development (the central business district's enumeration areas have the development concentrated in the low-lying areas, and the uniform-distribution assumption overestimates the exposure in these areas); (2) the 2020 census data is four years old, and the population and the buildings may have changed since the census (the metropolitan area's population has grown by approximately 2.5 per cent between 2020 and 2024, and the buildings have grown by

approximately 1.5 per cent). The issues are documented in the assessment's limitations section, and the recommendations for the future assessments are developed (the use of the building footprints from the high-resolution satellite imagery for the more accurate exposure estimation, and the use of the intercensal population estimates for the more current population data).

The joint authorization is conducted on 15 March 2024. The integrated product is authorized by the office's office manager (for the office's contribution) and the statistics agency's director (for the agency's contribution), with the authorization documenting the product's scope, the methodology, the limitations, and the classification. The product carries the Public classification for the general version (the population and the building exposure at the metropolitan-area and the neighbourhood level, without the enumeration-area-level detail) and the Internal classification for the operational version (the population and the building exposure at the enumeration-area level, with the detailed critical infrastructure information).

The joint publication is conducted on 25 March 2024. The integrated product is published on the office's website and the statistics agency's website, with the joint authorship ("Produced by the Geospatial Office and the National Statistics Agency"), the joint copyright, and the joint disclaimer. The Public version is published for the free access, with the open data licence that allows the free use and the redistribution. The Internal version is delivered to the emergency management authority, the planning ministry, and the other authorized recipients through the secure distribution channel. The publication is announced through the joint press release, with the office manager and the statistics agency's director jointly presenting the assessment at the press conference.

The assessment's use is demonstrated three weeks after the publication, when the March 2024 flood event (treated in Chapter 24, section 24.9) inundates approximately 18 square kilometres of the Portsmith metropolitan area. The emergency management authority uses the assessment's operational version for the evacuation planning (the identification of the 38,420 residents exposed to the 100-year flood hazard, the prioritization of the evacuation for the residents in the high-exposure enumeration areas), the resource positioning (the pre-positioning of the swift-water rescue teams in the high-exposure neighbourhoods), and the public education (the publication of the assessment's general version for the public awareness). During the flood event, the authority uses the assessment's operational version for the evacuation orders (the 6 evacuation zones, treated in Chapter 24, section 24.9, are based on the assessment's enumeration-area-level exposure), the resource deployment, and the public warnings. After the flood event, the authority uses the assessment's operational version for the post-event assessment (the comparison of the actual inundation with the 100-year flood hazard map, the comparison of the actual impact with the exposure estimates).

The post-event review of the assessment is conducted on 15 June 2024, jointly by the office and the statistics agency. The review covers the assessment's accuracy (the actual population affected by the March 2024 flood was 38,000, compared to the assessment's estimate of 38,420 for the 100-year flood hazard, an accuracy of 99 per cent for the population estimate), the assessment's usefulness (the assessment was used by the authority for the evacuation planning, the evacuation orders, and the post-event assessment, confirming the assessment's operational value), and the lessons learned (the area-weighted aggregation's assumption was adequate for the March 2024 event because the event's inundation extent matched the 100-year flood hazard extent, but the assumption may not be adequate for the events with the different inundation patterns; the 2020 census data was adequate for the March

2024 event because the population change between 2020 and 2024 was small, but the intercensal estimates should be used for the future assessments to maintain the accuracy). The review's findings are documented, and the documentation is the basis for the improvement of the assessment's methodology and the data for the next iteration.

This worked example illustrates the interface's full operation: the joint planning on 5 February 2024 (the office manager and the statistics agency's data liaison), the spatial component production on 12-16 February (the hazard specialist, the 1,247 enumeration areas, the inundation statistics), the statistical component production on 19-23 February (the methodological director, the 2020 census data, the confidentiality protection), the joint integration on 1-5 March (the geospatial join, the area-weighted aggregation, the 38,420 exposed population), the joint review on 8-12 March (the technical lead and the quality assurance team, the two issues identified), the joint authorization on 15 March (the office manager and the director), the joint publication on 25 March (the Public and the Internal versions, the joint press release), the assessment's use during the March 2024 flood event (the evacuation planning, the evacuation orders, the post-event assessment), and the post-event review on 15 June 2024 (the 99-per-cent accuracy, the lessons learned). The interface's data-sharing arrangement was maintained throughout: the office provided the spatial component, the statistics agency provided the statistical component, and the two components were integrated jointly. The interface's independence protection was maintained: the statistics agency's data was used for the office's exposure assessment, but the agency's independence was not compromised, and the agency's data was not used for the policy advocacy. The interface's confidentiality protection was maintained: the de-identification, the aggregation, and the small-cell suppression were applied, and the assessment's products did not reveal the individual-level data. The interface is the foundation of the office's statistical grounding, and the interface's integrity is the condition of the office's contribution to the jurisdiction's evidence-based governance. The March 2024 flood event demonstrated the interface's value: the assessment's exposure estimates, produced jointly by the office and the statistics agency, were used by the emergency management authority for the evacuation planning and the evacuation orders, and the assessment's 99-per-cent accuracy confirmed the interface's methodological soundness and the data-sharing arrangement's effectiveness. The interface is the mechanism that connects the geospatial office's spatial intelligence to the national statistics agency's statistical record, and the connection is the reason the interface exists.

Chapter 29 — The Geospatial Information System Specification

Part VI — Methods and Tools: The information system specification, the role family and qualifications, training, and continuous improvement.

29.1 The Architectural Pattern

The geospatial information system specification comprises six submodules — the spatial data infrastructure, the acquisition desk, the processing engine, the modeling workbench, the product studio, and the catalog and archive — integrated with the government's core administrative information system. The specification is examined as an architectural pattern, not as a product recommendation: the six submodules correspond to the six pipeline capabilities treated in Part II (Chapters 6 through 12), and the integration with the core system ensures that the geospatial work participates in the government's standard records, finance, human resources, procurement, and workflow infrastructure. The architectural pattern is the pattern that the geospatial office's information system should follow, regardless of the specific software products the office selects to implement the pattern.

The architectural pattern is derived from the pipeline architecture, treated in Chapter 6, and the pattern is the technical expression of the pipeline's shared-backbone principle. The six submodules correspond to the six capabilities: the spatial data infrastructure corresponds to the base layer capability (Chapter 8), the acquisition desk corresponds to the acquisition capability (Chapter 7), the processing engine corresponds to the processing capability (Chapter 9), the modeling workbench corresponds to the modeling capability (Chapter 10), the product studio corresponds to the production capability (Chapter 11), and the catalog and archive corresponds to the archive capability (Chapter 12). The correspondence is not coincidental: the information system is the technical implementation of the pipeline, and the pipeline is the doctrinal specification of the information system. An office that adopts the pipeline doctrine must adopt the information system architecture, and an office that adopts the information system architecture must adopt the pipeline doctrine.

The architectural pattern's defining property is the separation between the shared backbone and the swappable heads. The shared backbone comprises the spatial data infrastructure, the acquisition desk, the processing engine, the product studio, and the catalog and archive — the five submodules that are the same across the four domains (the flood risk, the food security, the energy infrastructure, and the minerals). The swappable heads comprise the modeling workbench's domain plug-ins — the hydrology plug-in, the crop model plug-in, the change detection plug-in, and the prospectivity plug-in — the four plug-ins that vary by domain. The separation allows the office to maintain a single shared stack for the five backbone submodules (with the consistent investment, the consistent maintenance, and the consistent integration), while serving the four domains through the four swappable plug-ins (with the domain-specific investment, the domain-specific maintenance, and the domain-specific validation). The separation is the architectural decision that makes the small office's multi-domain capability possible, and the separation is what distinguishes the geospatial information system from the single-purpose geospatial applications.

The architectural pattern is technology-agnostic. The office may implement the six submodules using the commercial off-the-shelf software (the ESRI ArcGIS platform, the Hexagon Geospatial

platform, the open-source geospatial stack with the PostGIS, the GeoServer, the QGIS, and the GDAL), or the office may implement the submodules using the custom-developed software. The architectural pattern does not specify the software; the architectural pattern specifies the capabilities that the software must provide. The office that selects the commercial software purchases the capabilities from the vendor; the office that selects the open-source software assembles the capabilities from the open-source components; the office that selects the custom development builds the capabilities from the internal resources. The three approaches have different cost, risk, and capability profiles, and the office's selection is the strategic decision that the office's leadership makes, with the technical lead's technical advice and the office manager's operational advice. The architectural pattern is the same regardless of the approach, and the pattern is the subject of this chapter.

29.2 The Spatial Data Infrastructure

The spatial data infrastructure is the submodule that maintains the office's base layers: the digital elevation model, the land cover, the hydrography, the administrative boundaries, and the cadastre feed. The submodule is the technical implementation of the base layer capability, treated in Chapter 8, and it provides the geographic context for every analysis the office conducts. The submodule comprises the base layer storage, the base layer versioning, the base layer synchronization, and the base layer access interface. The storage maintains the base layer data; the versioning maintains the base layer history; the synchronization maintains the base layer currency; and the access interface provides the base layer data to the other submodules and the specialists.

The base layer storage maintains the base layer data in the geospatial database, with the database configured for the spatial data's specific requirements. The geospatial database is typically a PostGIS database (the open-source spatial extension of the PostgreSQL database), with the PostGIS providing the spatial indexing (the GiST indexes for the spatial queries), the spatial functions (the ST_Intersects, the ST_Contains, the ST_Distance for the spatial analysis), and the spatial data types (the geometry, the geography, the raster for the spatial data). The database is configured with the appropriate storage parameters (the table partitioning for the large datasets, the cluster indexing for the frequently queried columns, the materialized views for the pre-computed aggregations), and the database is monitored for the performance (the query execution times, the index usage, the storage consumption). The base layer storage is maintained by the data pipeline operator, with the technical lead configuring the database parameters.

The base layer versioning maintains the base layer history, allowing the office to retrieve the base layer's state at any point in the past. The versioning is implemented through the versioned tables (each base layer record has the valid-from and the valid-to timestamps, with the records' history maintained in the same table or in the history table), and the versioning is queried through the temporal queries (the AS OF query retrieves the base layer's state at a specific time, the BETWEEN query retrieves the base layer's state over a time range). The versioning is the mechanism that allows the office to answer the audit questions ("What did the office know about the flood risk in district X on date Y?", treated in Chapter 12, section 12.3), and the versioning is the condition of the office's auditability.

The base layer synchronization maintains the base layer currency, ensuring that the office's copy of the base layer is synchronized with the authoritative source. The synchronization is conducted through the incremental updates (the authoritative source provides the changes since the last synchronization, and the office applies the changes to its copy), with the synchronization cadence determined by the base layer's update frequency (the daily for the cadastre feed from the cadastral authority, the annual for the land cover from the office's classification, the as-needed for the digital elevation model from the LiDAR acquisitions). The synchronization is monitored for the completeness and the timeliness, and the failures are resolved through the inter-agency coordination with the authoritative source.

The base layer access interface provides the base layer data to the other submodules and the specialists, with the interface implemented through the standard web services (the OGC Web Map Service for the rendered maps, the OGC Web Feature Service for the vector data, the OGC Web Coverage Service for the raster data) and the database connections (the direct database connections for the high-performance access). The interface is monitored for the availability and the performance, and the interface's documentation is maintained in the office's technical documentation. The access interface is the mechanism that makes the base layers available to the office's work, and the interface's reliability is the condition of the office's operational continuity.

The spatial data infrastructure is the foundation of the information system, because every other submodule depends on the base layers. The processing engine uses the base layers to georeference the acquired data; the modeling workbench uses the base layers to provide the geographic context for the analysis; the product studio uses the base layers to provide the cartographic context for the products; and the catalog and archive uses the base layers to link the products to their geographic context. The spatial data infrastructure that fails undermines every other submodule, and the spatial data infrastructure's integrity is the non-negotiable precondition for the information system's integrity. The submodule is maintained by the data pipeline operator, with the technical lead configuring the maintenance parameters, and the submodule's integrity is verified through the weekly base layer integrity audit, treated in the job breakdown sheet in section 15.5.

29.3 The Acquisition Desk

The acquisition desk is the submodule that manages the satellite tasking, the imagery purchasing, the sensor feed ingestion, and the data license management. The submodule is the technical implementation of the acquisition capability, treated in Chapter 7, and it is the entry point of the pipeline. The submodule comprises the tasking request interface, the purchase order interface, the sensor feed subscription interface, the license record management, and the ingestion verification.

The tasking request interface allows the specialists to submit the satellite tasking requests, with the requests specifying the area of interest (the polygon in the GeoJSON format), the acquisition window (the start and the end dates), the sensor parameters (the resolution, the bands, the incidence angle), and the priority (the routine, the urgent, the emergency). The interface validates the request against the tasking rules (the area of interest must be within the office's jurisdiction, the acquisition window must be in the future, the sensor parameters must be within the satellite's capabilities), and the interface submits the validated request to the satellite operator's tasking API (the Maxar API, the Airbus API, the Planet API). The interface tracks the request's status (the submitted, the accepted, the

scheduled, the acquired, the delivered, the failed), and the interface notifies the specialist when the tasking is delivered or the failure occurs. The tasking request interface is operated by the specialists, with the office manager authorizing the emergency tasking.

The purchase order interface allows the data pipeline operator to submit the imagery purchase orders, with the orders specifying the imagery product (the archive imagery or the newly-tasked imagery), the product specifications (the resolution, the bands, the processing level), the quantity (the square kilometres or the scenes), and the cost (the unit price and the total price). The interface validates the order against the budget (the order's total cost must be within the available budget), the procurement rules (the order must be under the framework agreement or the competitive procurement), and the license terms (the order's license must be compatible with the office's intended use). The interface submits the validated order to the satellite operator's purchase API, and the interface tracks the order's status (the submitted, the confirmed, the delivered, the invoiced, the paid). The purchase order interface is operated by the data pipeline operator, with the office manager authorizing the purchases above the threshold (typically $10,000).

The sensor feed subscription interface manages the sensor feed subscriptions, with the subscriptions specifying the sensor (the Sentinel-1, the Sentinel-2, the Sentinel-3, the Landsat 8, the Landsat 9), the area of interest (the office's jurisdiction), the acquisition criteria (the cloud cover threshold, the incidence angle range), and the delivery channel (the direct broadcast, the data hub, the cloud storage). The interface manages the subscriptions' lifecycle (the creation, the modification, the suspension, the termination), and the interface monitors the subscriptions' performance (the delivery timeliness, the delivery completeness, the data quality). The sensor feed subscription interface is operated by the data pipeline operator, with the technical lead configuring the subscription parameters.

The license record management records the data license terms for every acquired dataset, with the records specifying the license type (the open license, the commercial license, the restricted license), the license terms (the permitted uses, the prohibited uses, the redistribution rights, the attribution requirements), the license period (the start and the end dates), and the license cost (the one-time cost or the recurring cost). The records are maintained in the catalog and archive submodule, and the acquisition desk provides the records to the catalog and archive at the moment of acquisition. The license records are enforced at the production and archive stages, where the products are checked against the license terms before release. The license record management is the legal integrity guarantee of the pipeline, and the management is the mechanism that prevents the office from releasing products that infringe the data providers' intellectual property.

The ingestion verification verifies that the acquired data was ingested correctly into the processing engine. The verification checks the data's completeness (the expected files are present), the data's format (the files are in the expected format), the data's georeferencing (the data is georeferenced to the expected coordinate system), and the data's quality (the data is within the expected quality parameters, such as the cloud cover, the geometric accuracy, the radiometric accuracy). The verification's results are recorded in the acquisition desk's log, and the failures are resolved through the inter-agency coordination with the data provider. The ingestion verification is the mechanism that ensures the acquired data is ready for the processing, and the verification is the condition of the processing engine's effectiveness.

29.4 The Processing Engine

The processing engine is the submodule that performs the georeferencing, the calibration, the mosaicking, the projection management, the cloud masking, and the data fusion. The submodule is the technical implementation of the processing capability, treated in Chapter 9, and it is the submodule that transforms the raw acquired data into the analysis-ready products. The submodule comprises the processing queue management, the processing parameter management, the ground control point library, the processing output catalog, and the validation interface.

The processing queue management manages the processing jobs, with the jobs submitted by the data pipeline operator or the specialists. The queue management prioritizes the jobs (the emergency processing for the event-driven response, the urgent processing for the time-sensitive analysis, the routine processing for the standing production), allocates the computing resources (the CPU, the GPU, the memory, the storage), and tracks the jobs' status (the queued, the running, the completed, the failed). The queue management is implemented through the workflow engine (the Apache Airflow, the Pegasus, the custom-developed workflow engine), with the engine managing the job dependencies (the processing stages that must be completed before the next stage), the job retries (the failed jobs are retried with the specified retry policy), and the job notifications (the operator and the specialist are notified when the job completes or fails).

The processing parameter management manages the processing parameters, with the parameters specified for each processing stage and each sensor. The parameters include the georeferencing parameters (the polynomial order, the ground control point selection, the transformation type), the calibration parameters (the atmospheric model, the aerosol optical thickness, the water vapour, the ozone), the mosaicking parameters (the colour balancing method, the seamline generation, the feathering distance), the projection parameters (the source projection, the target projection, the transformation method), the cloud masking parameters (the spectral thresholds, the geometric tests, the temporal tests), and the data fusion parameters (the fusion method, the band selection, the resolution ratio). The parameters are stored in the parameter library, with the library versioned and the changes logged. The technical lead configures the parameters, and the data pipeline operator applies the parameters in the processing runs.

The ground control point library maintains the ground control points used for the georeferencing, with the points sourced from the high-accuracy reference data (the GPS-surveyed points, the high-resolution orthophotos, the existing georeferenced imagery). The library stores the points' coordinates (the latitude, the longitude, the elevation), the points' accuracy (the horizontal and the vertical accuracy), the points' description (the feature type, the location description), and the points' source (the survey, the orthophoto, the imagery). The library is maintained by the technical lead, with the points added, modified, and retired as the reference data is updated. The ground control point library is the geometric integrity foundation of the processing engine, and the library's quality is the condition of the georeferencing's accuracy.

The processing output catalog catalogs the processed products, with the catalog records specifying the product's identifier, the product's type (the orthorectified imagery, the surface reflectance, the cloud-masked imagery, the fused product), the product's source (the acquisition desk's identifier), the product's processing parameters (the parameter library's version), the product's quality (the geometric accuracy, the radiometric accuracy, the cloud cover), and the product's storage

location (the file path or the database record). The catalog is the mechanism that makes the processed products retrievable, and the catalog is the input to the modeling workbench and the product studio.

The validation interface allows the technical lead to validate the processed products, with the validation comparing the products' geometric accuracy against the independent check points (the points not used in the georeferencing), the products' radiometric accuracy against the reference targets (the ground targets with the known reflectance), and the products' completeness against the expected coverage. The validation's results are recorded in the catalog and archive submodule, and the failures are resolved through the parameter adjustments or the reprocessing. The validation interface is the geometric and radiometric integrity guarantee of the processing engine, and the validation is the condition of the processed products' analysis-readiness.

29.5 The Modeling Workbench

The modeling workbench is the submodule that hosts the domain-specific analytical models, and it is the submodule that implements the swappable-heads design, treated in Chapter 6. The submodule comprises the workbench environment, the domain plug-ins, the parameter management, the validation framework, and the provenance recording.

The workbench environment provides the shared environment for the domain plug-ins, with the environment comprising the data access interface (the interface to the spatial data infrastructure's base layers and the processing engine's outputs), the computing resources (the CPU, the GPU, the memory, the storage allocated to the model runs), the model execution engine (the engine that executes the plug-ins' code, with the engine supporting the Python, the R, the Fortran, and the C++ languages), and the model output interface (the interface that delivers the model outputs to the product studio and the catalog and archive). The workbench environment is the same across the four domains, and the environment is the shared stack that the four plug-ins run on.

The domain plug-ins implement the domain-specific methods. The hydrology plug-in (treated in Chapter 10, section 10.2) implements the rainfall-runoff modelling (the HEC-HMS model, the GR4J model), the hydraulic modelling (the HEC-RAS 2D model, the Delft3D model), and the inundation mapping (the post-processing of the hydraulic model's outputs into the inundation extents and the depths). The crop model plug-in (treated in Chapter 10, section 10.3) implements the vegetation index computation (the NDVI, the EVI, the NDWI), the phenological stage determination (the accumulated growing degree days), the yield forecasting (the statistical regression, the mechanistic models such as the WOFOST and the DSSAT), and the supply-demand analysis. The change detection plug-in (treated in Chapter 10, section 10.4) implements the image differencing, the object-based change detection, the encroachment classification, and the vulnerability scoring. The prospectivity plug-in (treated in Chapter 10, section 10.5) implements the knowledge-driven prospectivity (the boolean logic, the fuzzy logic, the weights of evidence), the data-driven prospectivity (the logistic regression, the random forest, the neural networks), and the resource estimation (the kriging, the inverse distance weighting).

The parameter management manages the model parameters, with the parameters specified for each plug-in and each application. The parameters include the hydrological parameters (the curve numbers, the Manning's n values, the rainfall-runoff parameters), the crop parameters (the crop coefficients, the phenological thresholds, the regression coefficients), the change detection

parameters (the change thresholds, the buffer widths, the vulnerability scoring weights), and the prospectivity parameters (the weights, the evidence layers, the estimation parameters). The parameters are stored in the parameter library, with the library versioned and the changes logged. The technical lead configures the parameters, and the specialists apply the parameters in the model runs.

The validation framework provides the validation tools for the model runs, with the tools comparing the model's predictions against the observed data. The hydrology validation tools compare the model's predicted river stages with the observed stages at the gauge locations, the model's predicted inundation extents with the observed extents, and the model's predicted flood durations with the observed durations. The crop model validation tools compare the model's predicted yields with the reported yields, the model's predicted conditions with the observed conditions, and the model's predicted phenological stages with the observed stages. The change detection validation tools compare the model's detected changes with the field-verified changes, the model's vulnerability scores with the actual failure occurrences, and the model's encroachment classifications with the field classifications. The prospectivity validation tools compare the model's predicted prospectivity with the actual mineral occurrences, the model's resource estimates with the drilling results, and the model's prediction confidence with the actual discovery rates. The validation's results are recorded in the provenance recording, and the results are the basis for the model's improvement.

The provenance recording records the model version, the parameters, the run date, the validation results, and the output storage location, with the recording linked to the catalog and archive submodule. The provenance recording is the mechanism that makes the model runs reproducible (the same model version, the same parameters, the same input data produce the same output), auditable (the model's predictions can be traced to the specific model version and parameters), and defensible (the model's predictions can be justified with the provenance record). The provenance recording is the analytical integrity guarantee of the modeling workbench, and the recording is the condition of the model's credibility.

29.6 The Product Studio

The product studio is the submodule that performs the cartographic design, the symbology application, the classification marking, and the format conversion. The submodule is the technical implementation of the production capability, treated in Chapter 11, and it is the submodule that transforms the model outputs into the final deliverables. The submodule comprises the cartographic template management, the symbology library, the classification marking tools, and the format conversion tools.

The cartographic template management manages the cartographic templates, with the templates specifying the product's layout (the map area, the legend, the title block, the classification marking, the metadata block), the product's scale (the 1:25,000, the 1:100,000, the 1:500,000, with each scale having the generalization rules), the product's orientation (the portrait, the landscape, the custom), and the product's size (the A4, the A3, the A1, the custom). The templates are stored in the template library, with the library versioned and the changes logged. The production operator selects the appropriate template for each product, and the template is applied automatically by the product studio. The cartographic template management is the mechanism that ensures the product's visual consistency, and the management is the condition of the product's professionalism.

The symbology library maintains the symbology for the product's features, with the symbology specifying the feature's appearance (the colour, the line style, the fill pattern, the symbol, the label). The symbology is organized by the domain (the flood symbology with the blue gradient for the depth, the food security symbology with the green-to-red gradient for the condition, the energy infrastructure symbology with the red for the new features and the blue for the removed features, the minerals symbology with the warm-colour gradient for the high prospectivity), and the symbology is organized by the feature type (the points, the lines, the polygons, the rasters). The symbology library is maintained by the production operator, with the technical lead reviewing the symbology for the analytical accuracy. The symbology library is the mechanism that ensures the product's analytical clarity, and the library is the condition of the product's communicative effectiveness.

The classification marking tools apply the classification markings to the products, with the markings specified for each classification level (the Public, the Internal, the Confidential, the Security Confidential). The markings are applied in the standard format (the capital letters, the standard font, the standard size) and in the standard locations (the top and the bottom of every page, on the cover, in the metadata). The tools verify the markings' completeness before the product's release, with the verification checking that the markings are present in all the required locations and that the markings are consistent with the product's classification. The classification marking tools are the security integrity guarantee of the production function, and the tools are the mechanism that prevents the release of the incorrectly classified products.

The format conversion tools convert the products to the consumer-required formats, with the formats including the GeoTIFF (for the raster products, preserving the georeferencing), the Shapefile (for the vector products, preserving the symbology through the companion files), the GeoPackage (for the vector products, with the single-file format), the KML (for the products viewed in the virtual globe viewers), the PDF (for the products read by the decision-makers), and the DOCX (for the report products). The format conversion tools are configured with the format-specific parameters (the GeoTIFF's compression, the Shapefile's coordinate system, the KML's scale-dependent rendering, the PDF's resolution), and the tools verify the conversion's completeness (the converted product contains all the original product's information, with the georeferencing, the symbology, and the classification marking preserved).

The product studio is maintained by the production operator, who produces the products, and the office manager, who reviews the products before release. The submodule's discipline is the foundation of the pipeline's communicative and security integrity, and the discipline ensures the products are clear, correctly classified, and in the right format. The product studio's effectiveness is verified through the production operator's classification marking verification (treated in the job breakdown sheet in section 16.4) and the office manager's pre-release review.

29.7 The Catalog and Archive

The catalog and archive is the submodule that preserves the office's outputs and maintains the catalog that makes them retrievable. The submodule is the technical implementation of the archive capability, treated in Chapter 12, and it is the submodule that maintains the complete provenance chain from acquisition through archive. The submodule comprises the product catalog, the version

history, the classification markings, the retention designations, the access logging, and the provenance chain.

The product catalog indexes the products, with the catalog records specifying the product's identifier (the unique identifier, typically the UUID), the product's type (the map, the bulletin, the advisory, the alert, the report), the product's title, the product's author (the specialist who produced the product), the product's date (the production date and the release date), the product's classification (the Public, the Internal, the Confidential, the Security Confidential), the product's geographic coverage (the bounding box or the polygon), the product's temporal coverage (the start and the end dates), the product's storage location (the file path or the database record), and the product's access permissions (the authorized users and roles). The catalog is the primary retrieval interface for the products, and the catalog's completeness and accuracy are the conditions of the products' retrievability.

The version history maintains the product versions, with the versions created when the product is updated (the product's content is changed, the product's classification is changed, the product's format is changed). The version history records the version's identifier, the version's date, the version's author, the version's changes (the description of what was changed), and the version's storage location. The version history is the mechanism that allows the office to retrieve the product's state at any point in the past, and the version history is the condition of the office's auditability (the audit's question "What did the office know about the flood risk in district X on date Y?" is answered by retrieving the product's version current on that date).

The classification markings record the product's classification, with the markings applied by the production operator and verified by the office manager. The markings are stored in the catalog's record, and the markings are enforced by the access control (the access control restricts the product's access to the authorized users, with the authorization determined by the product's classification and the user's clearance). The classification markings are the security integrity guarantee of the catalog and archive, and the markings are the mechanism that prevents the unauthorized access to the classified products.

The retention designations manage the product's lifecycle, with the designations specifying the product's retention period (the minimum period the product must be retained) and the product's disposition (the action taken at the end of the retention period: the permanent retention, the archival transfer, the destruction). The retention periods are specified by the product's type and the office's retention schedule (the crop condition bulletins retained for 10 years, the infrastructure risk maps retained for 7 years aligned to the regulator's compliance cycle, the commercial minerals commissions retained for 10 years after the completion for the contractual disputes, treated in Chapter 12, section 12.4). The retention designations are the mechanism that ensures the office's compliance with the records management requirements, and the designations are the condition of the office's legal defensibility.

The access logging records the access to the classified products, with the logging specifying the user, the date, the time, the product, and the action (the view, the download, the modification). The access logging is the mechanism that allows the office to audit the classified products' access, and the logging is the condition of the classified products' security. The access logging is reviewed by the production operator in the weekly classified access log review (treated in the job breakdown sheet in section 16.6) and by the office manager in the monthly compliance review.

The provenance chain links the products to their source data, processing, modeling, and production, with the chain recording the product's complete history from the acquisition through the archive. The chain records the source data (the acquisition desk's identifier, the sensor, the acquisition date), the processing (the processing engine's identifier, the processing parameters, the processing date), the modeling (the modeling workbench's identifier, the model version, the parameters, the run date), the production (the product studio's identifier, the template, the symbology, the classification marking, the production date), and the archive (the catalog and archive's identifier, the catalog record, the version, the storage location). The provenance chain is the integrity guarantee of the pipeline, and the chain makes the office's products auditable, reproducible, and defensible.

29.8 The Integration with the Core System

The integration with the government's core administrative information system is the integration that ensures the geospatial work participates in the government's standard records, finance, human resources, procurement, and workflow infrastructure. The core system comprises the records management, the financial management, the human resources, the procurement, and the workflow management. The geospatial information system integrates with the core system through each of these components, and the integration is the mechanism that makes the geospatial office's work a part of the government's standard operations.

The records management integration links the catalog and archive submodule to the core records management, with the integration ensuring that the office's products are managed under the government's records management policy. The integration comprises the records schedule (the office's retention schedule, treated in Chapter 12, section 12.4, is aligned with the government's records schedule), the records transfer (the products that are designated for the archival transfer are transferred to the government's archives at the end of the retention period), and the records disposal (the products that are designated for the destruction are destroyed in accordance with the government's disposal procedures, with the destruction documented in the disposal log). The records management integration is the mechanism that ensures the office's compliance with the government's records management requirements, and the integration is the condition of the office's legal defensibility.

The financial management integration links the acquisition desk and the commercial revenue to the core financial management, with the integration ensuring that the office's expenditures and revenues are managed under the government's financial management framework. The integration comprises the purchase order integration (the acquisition desk's purchase orders are recorded in the core financial management, with the expenditures tracked against the budget), the commercial revenue integration (the commercial minerals commissions' revenue is recorded in the core financial management, with the revenue tracked against the commercial revenue account, treated in Chapter 5, section 5.3), and the cost accounting integration (the office's costs are allocated to the public-service appropriation or the commercial charge codes, with the allocation tracked in the core financial management). The financial management integration is the mechanism that ensures the office's compliance with the government's financial management requirements, and the integration is the condition of the commercial firewall's effectiveness (treated in Chapter 5).

The human resources integration links the role family (treated in Chapter 4) to the core human resources, with the integration ensuring that the office's personnel are managed under the government's human resources framework. The integration comprises the position management (the office's positions are defined in the core human resources, with the positions' classifications, the reporting lines, and the responsibilities documented), the personnel management (the office's personnel are managed in the core human resources, with the personnel's assignments, the performance reviews, and the training records tracked), and the security clearance management (the office's personnel's security clearances are managed in the core human resources, with the clearances' levels, the granting dates, and the review dates tracked). The human resources integration is the mechanism that ensures the office's compliance with the government's human resources requirements, and the integration is the condition of the office's personnel management.

The procurement integration links the acquisition desk to the core procurement, with the integration ensuring that the office's procurement is managed under the government's procurement framework. The integration comprises the framework agreement management (the office's framework agreements with the satellite operators and the commercial service providers are recorded in the core procurement, with the agreements' terms, the value, and the expiry dates tracked), the competitive procurement management (the office's competitive procurements are conducted through the core procurement, with the tendering, the evaluation, and the award documented), and the procurement compliance (the office's procurements are checked against the procurement rules, with the rules' compliance verified before the purchase order's approval). The procurement integration is the mechanism that ensures the office's compliance with the government's procurement requirements, and the integration is the condition of the office's procurement's legality.

The workflow management integration links all submodules to the core workflow, with the integration ensuring that the office's tasks are managed under the government's workflow framework. The integration comprises the task management (the office's tasks are created, assigned, tracked, and completed in the core workflow, with the tasks linked to the products and the projects), the approval management (the office's approvals are managed in the core workflow, with the approvals' routing, the approvers, and the decisions tracked), and the notification management (the office's notifications are managed in the core workflow, with the notifications' triggers, the recipients, and the delivery tracked). The workflow management integration is the mechanism that ensures the office's compliance with the government's workflow requirements, and the integration is the condition of the office's operational efficiency.

29.9 Worked Example: The System Specification for the Riverbourne Basin Flood Event Response

The system specification is illustrated through the Riverbourne Basin flood event of March 2024, treated in Chapters 24, 27, and 28. The event's response required the geospatial information system's six submodules to operate in coordination, with each submodule contributing its specific capability to the event-driven response. The worked example traces the event's timeline through the six submodules, demonstrating how the architectural pattern's separation between the shared backbone and the swappable heads enabled the office's effective response.

The event begins on 23 March 2024 at 09:15, when the meteorological service issues the severe weather warning. The data pipeline operator, monitoring the meteorological feeds through the acquisition desk's sensor feed subscription interface, notes the warning and notifies the office manager. The acquisition desk's sensor feed subscription interface had been configured to subscribe to the meteorological service's warning feed, with the subscription delivering the warnings to the acquisition desk's notification system. The acquisition desk's role in the event's beginning is the detection of the triggering event, and the submodule's sensor feed subscription interface is the mechanism that enabled the detection.

The office manager activates the office's event-driven response at 09:45 (treated in Chapter 24, section 24.9). The activation triggers the modeling workbench's hydrology plug-in, with the plug-in configured for the event-specific run. The modeling workbench's workbench environment provides the data access interface to the spatial data infrastructure's base layers (the LiDAR DEM at 2-metre resolution, treated in Chapter 19, section 19.9) and the processing engine's outputs (the most recent Sentinel-1 and Sentinel-2 imagery, processed and validated). The workbench environment's data access interface is the mechanism that provided the hydrology plug-in with the input data, and the interface's reliability is the condition of the plug-in's effectiveness.

The hazard specialist configures the hydrology plug-in at 10:00 (treated in Chapter 27, section 27.9). The HEC-HMS model is configured with the observed rainfall (from the rain gauge network, accessed through the spatial data infrastructure's hydrography base layer) and the forecast rainfall (from the meteorological service's numerical weather prediction model, accessed through the acquisition desk's sensor feed subscription interface). The HEC-RAS 2D model is configured with the LiDAR DEM, the channel geometry (from the water authority's HEC-RAS geometry files, accessed through the acquisition desk's data-sharing interface), the levee data (from the water authority's Shapefiles, with the Levee Section 23's crest elevation adjusted for the 0.15-metre settlement, treated in Chapter 27, section 27.9), and the boundary conditions. The modeling workbench's parameter management stores the event-specific parameters, and the parameter library's version is recorded in the provenance recording.

The HEC-RAS run is executed at 11:00, with the modeling workbench's model execution engine allocating the computing resources (the 16 CPU cores, the 64 GB of memory, the 2 TB of storage for the model's temporary files) and the run taking approximately 90 minutes. The model execution engine's workflow management handles the model's dependencies (the HEC-HMS run must complete before the HEC-RAS run can begin) and the model's notifications (the hazard specialist is notified when the run completes). The modeling workbench's validation framework validates the model's outputs (the predicted river stages are compared with the observed stages at the gauge locations, with the validation's results recorded in the provenance recording).

The hazard specialist produces the event-specific flood inundation map at 12:30, with the map showing the no-breach and the breach scenarios (treated in Chapter 27, section 27.9). The map is produced in the product studio, with the cartographic template management applying the flood event template (the standard template for the flood inundation maps, with the blue gradient for the depth, the red for the evacuation zones, the yellow for the shelters), the symbology library applying the flood symbology, the classification marking tools applying the Internal classification (with the Security Confidential version produced separately for the critical infrastructure information), and the format

conversion tools producing the GeoTIFF (for the ingestion into the emergency management authority's common operating picture platform), the Shapefile (for the evacuation zones), the KML (for the field teams' virtual globe viewers), and the PDF (for the decision-makers' reading).

The production operator applies the classification markings and produces the products at 13:00, with the catalog and archive submodule recording the products in the product catalog. The catalog records specify the products' identifiers (the UUIDs), the products' types (the flood inundation map, the Internal and the Security Confidential versions), the products' authors (the hazard specialist), the products' dates (the production date 23 March 2024, the release date 23 March 2024), the products' classification (the Internal and the Security Confidential), the products' geographic coverage (the Portsmith metropolitan area, 18 square kilometres), the products' temporal coverage (the forecast horizon 24-25 March 2024), the products' storage locations (the file paths in the office's classified archive), and the products' access permissions (the authorized users: the office manager, the hazard specialist, the production operator; the authorized recipients: the emergency management authority, the water authority).

The catalog and archive submodule also records the products' provenance chain, with the chain linking the products to their source data (the acquisition desk's identifiers for the Sentinel-1, the Sentinel-2, and the meteorological service's forecast), their processing (the processing engine's identifiers for the georeferenced and the calibrated imagery), their modeling (the modeling workbench's identifiers for the HEC-HMS and the HEC-RAS runs, with the model versions, the parameters, and the run dates), and their production (the product studio's identifiers for the cartographic templates, the symbology, and the classification markings). The provenance chain is the integrity guarantee of the event-specific products, and the chain is the mechanism that allows the office to defend the products' accuracy and the methodology in the post-event review.

The products are delivered to the emergency management authority at 13:30 (treated in Chapter 24, section 24.9) and to the water authority at 13:00 (treated in Chapter 27, section 27.9), with the delivery conducted through the secure data interface. The secure data interface is the spatial data infrastructure's access interface, configured for the inter-agency data sharing, with the interface operating over the HTTPS protocol with the mutual TLS authentication. The interface's reliability is the condition of the products' delivery, and the interface's logging is the mechanism that records the delivery for the audit trail.

The model integration with the water authority is conducted through the modeling workbench's data interface, with the office's HEC-RAS model exchanging the data with the water authority's operational models every 6 hours (treated in Chapter 27, section 27.9). The data interface uses the HEC-DSS format for the time series data and the GeoTIFF format for the raster data, with the interface's standardized formats ensuring the two offices' models' interoperability. The model integration is the mechanism that allows the two offices to produce a coherent flood risk picture, and the integration's effectiveness is the condition of the coordinated response.

The event-driven cadence continues through 23-26 March, with the modeling workbench producing the updated model runs every 6 hours, the product studio producing the updated products, the catalog and archive recording the updated versions, and the spatial data infrastructure's access interface delivering the products to the emergency management authority and the water authority. The event-driven cadence's intensity requires the system's scalability, with the modeling workbench's

computing resources scaled up (the additional CPU cores and the memory allocated for the accelerated model runs), the product studio's production capacity scaled up (the additional production operators assigned to the event-driven production), and the catalog and archive's storage capacity scaled up (the additional storage allocated for the event's products). The system's scalability allows the office to maintain the event-driven cadence through the event's duration, and the scalability is the condition of the office's event-driven response's effectiveness.

The event recedes on 26-28 March, and the office stands down the event-driven response on 28 March. The catalog and archive submodule's retention designations are applied to the event's products, with the products designated for the 10-year retention (the standard retention period for the flood inundation maps, treated in Chapter 12, section 12.4). The access logging records the access to the event's classified products, with the logging reviewed in the weekly classified access log review. The provenance chain is the basis for the post-event review's reconstruction of the event's timeline, and the chain is the basis for the model validation's comparison of the model's predictions with the observed data.

The post-event review is conducted on 5 April 2024 (treated in Chapter 24, section 24.9), and the post-event assessment is produced on 15 June 2024. The system's role in the event's response is assessed, with the assessment identifying the system's strengths (the sensor feed subscription interface's timely detection, the modeling workbench's effective event-specific configuration, the product studio's timely production, the catalog and archive's complete provenance chain) and the system's weaknesses (the model's storm surge boundary condition's underestimation, the levee settlement data's absence from the standing waterway data sharing). The assessment's recommendations include the improvement of the storm surge boundary condition (the use of the ensemble forecast rather than the deterministic forecast), the improvement of the waterway data sharing (the inclusion of the levee settlement data in the standing data sharing), and the improvement of the system's scalability (the additional computing resources for the future event-driven responses). The recommendations are tracked through the monthly coordination meetings, and the improvements are implemented before the 2025 flood season.

This worked example illustrates the system specification's full operation: the spatial data infrastructure's base layers and access interface (the LiDAR DEM, the hydrography, the secure data interface), the acquisition desk's sensor feed subscription and data-sharing interface (the meteorological warnings, the waterway data), the processing engine's outputs (the georeferenced and the calibrated imagery), the modeling workbench's hydrology plug-in (the HEC-HMS and the HEC-RAS runs, the 6-hour update cycle), the product studio's cartographic template and symbology (the flood event template, the blue gradient, the Internal and the Security Confidential markings), and the catalog and archive's product catalog, provenance chain, and retention designations (the 10-year retention, the access logging). The architectural pattern's separation between the shared backbone and the swappable heads enabled the office's effective response: the shared backbone (the spatial data infrastructure, the acquisition desk, the processing engine, the product studio, the catalog and archive) operated consistently across the event, while the swappable head (the hydrology plug-in) provided the domain-specific analytical capability. The system's integration with the core system (the records management, the financial management, the human resources, the procurement, the workflow management) ensured the office's compliance with the government's standard operations.

The system specification is the technical foundation of the office's geospatial intelligence capability, and the specification's integrity is the condition of the office's contribution to the public safety during the flood events. The Riverbourne Basin flood event demonstrated the system specification's value: the six submodules' coordination, the swappable-heads design's flexibility, and the provenance chain's auditability enabled the office's effective event-driven response, and the response's effectiveness is the reason the system specification exists.

Chapter 30 — The Geospatial Role Family and Qualifications

30.1 The Role Family

The geospatial role family defines six seats with two domain specializations, and it is the formal input to the qualifications authority for the establishment of the geospatial role family in the government's rank structure. The six seats are the office manager, the technical lead, the data pipeline operator, the production operator, the hazard specialist, and the resource economics specialist, treated in Part III (Chapters 13 through 18). The two domain specializations are the hazard specialization (the flood risk and energy infrastructure domains, treated in Chapters 19 and 21) and the resource economics specialization (the food security and minerals domains, treated in Chapters 20 and 22). The role family is distinct from the generalist administrative stream, because the work requires specialized technical knowledge that the generalist stream does not provide, and distinct from the IT professional stream, because the work requires analytical and cartographic judgment that the IT stream does not provide.

The role family is the institutional expression of the geospatial profession within the government's rank structure, and the role family is the mechanism by which the geospatial profession is established as a recognized career stream. The role family that is ratified by the qualifications authority becomes a permanent feature of the government's rank structure, with the ranks, the qualifications, the training, and the payscale that the role family specifies. The role family that is not ratified remains an informal arrangement, with the geospatial work conducted by generalist administrators or IT professionals who lack the specialized knowledge the work requires. The ratification is the condition of the role family's permanence, and this chapter is the formal submission for the ratification.

The role family's six seats are organized into three functional groups. The leadership group comprises the office manager and the technical lead, with the office manager responsible for the office's institutional relationships and the technical lead responsible for the office's technical standards. The pipeline group comprises the data pipeline operator and the production operator, with the data pipeline operator responsible for the data acquisition and processing and the production operator responsible for the product generation and archive. The specialist group comprises the hazard specialist and the resource economics specialist, with the hazard specialist responsible for the flood risk and energy infrastructure domains and the resource economics specialist responsible for the food security and minerals domains. The three groups are operationally interdependent, with the leadership group providing the direction, the pipeline group providing the data and the products, and the specialist group providing the analytical intelligence.

The role family's two domain specializations reflect the analytical depth that the four domains require. The hazard specialization covers the flood risk and energy infrastructure domains, which share the hazard-monitoring character: both domains monitor the physical conditions (the flood hazard, the infrastructure encroachment) that may produce the hazardous events, and both domains produce the intelligence that informs the hazard management. The resource economics specialization covers the food security and minerals domains, which share the resource-economics character: both domains assess the resources (the crop yield, the mineral prospectivity) that have the economic value, and both domains produce the intelligence that informs the resource management. The two

specializations are distinct because the analytical methods, the data sources, and the consumer relationships differ between the hazard and the resource economics domains, and the specialization ensures that the specialists have the depth of knowledge their domains require.

30.2 The Rank Structure

The rank structure of the geospatial role family follows the government's general rank structure, with an operational tier and an executive tier. The operational tier comprises the data pipeline operator, the production operator, the two specialists, and the technical lead, corresponding to the operational ranks of the general civil service. The executive tier comprises the office manager, corresponding to the executive ranks. Within the operational tier, the ranks are differentiated by experience and responsibility, with the operators at the entry-level operational ranks, the specialists at the senior operational ranks, and the technical lead at the bridge rank between the operational and executive tiers. The rank structure is the same as that described in Chapter 4, section 4.3, and this chapter constitutes the formal submission of the rank structure to the qualifications authority.

The operational tier's entry-level ranks are the ranks at which the operators and the junior specialists enter the role family. The data pipeline operator enters at the entry-level operational rank (the rank equivalent to the general civil service's entry-level technical rank), with the entry requirements comprising the technical qualification in geospatial technology and the basic certifications. The production operator enters at the same entry-level operational rank, with the entry requirements comprising the technical qualification in geospatial technology and the basic certifications. The junior specialist enters at the entry-level operational rank or the next-higher rank, depending on the academic qualification (the bachelor's degree enters at the entry-level rank, the master's degree enters at the next-higher rank), with the entry requirements comprising the university degree in the domain and the domain-specific certifications.

The operational tier's senior ranks are the ranks at which the senior specialists and the technical lead serve. The senior specialist serves at the senior operational rank (the rank equivalent to the general civil service's senior technical rank), with the promotion requirements comprising the years of experience (typically 5 to 8 years), the demonstrated competence in the domain, and the advanced certifications. The technical lead serves at the bridge rank (the rank equivalent to the general civil service's principal technical rank), with the promotion requirements comprising the years of experience (typically 8 to 12 years), the demonstrated leadership in the technical standards, and the leadership certifications. The bridge rank is the rank from which the officer may be promoted to the executive tier, and the bridge rank is the mechanism that allows the technical career to progress to the executive level without leaving the technical specialization.

The executive tier comprises the office manager, who serves at the executive rank (the rank equivalent to the general civil service's executive rank). The promotion to the executive rank requires the years of experience (typically 12 to 15 years), the demonstrated leadership in the office management, the management qualifications, and the selection through the competitive process. The executive rank is the rank at which the office manager leads the office, represents the office in the inter-agency coordination, and reports to the desk chief and the bureau chief. The executive rank is the highest rank in the geospatial role family, and the rank is the career's culmination for the officers who progress through the role family's ranks.

The rank structure is designed to be compatible with the government's general rank structure, so that geospatial officers can move between the geospatial role family and the generalist administrative stream or the IT professional stream. The compatibility is important because it ensures that the geospatial role family does not become a closed career, and that the officers who enter the role family can exit it if the office is dissolved or if their career aspirations change. The compatibility is also important because it allows the office to recruit from the generalist and IT streams, and to retain officers who may later move to other parts of the government. The rank structure is the institutional framework for the geospatial career, and the framework is the condition of the role family's viability.

30.3 The Qualifications

The qualifications for the geospatial role family comprise the academic qualifications, the technical certifications, and the professional certifications. The three categories are complementary, with the academic qualifications establishing the foundational knowledge, the technical certifications establishing the tool-specific competence, and the professional certifications establishing the professional standing. The qualifications are specified for each seat and each rank, and the qualifications are the conditions of the officer's appointment and promotion.

The academic qualifications are the university degrees or equivalent technical qualifications that establish the officer's foundational knowledge in geospatial science, geography, geology, hydrology, agriculture, or a related field. The operators require a technical qualification in geospatial technology or a related field, typically a diploma or an associate degree from a technical college or a polytechnic. The specialists require a university degree in their domain: the hazard specialist requires a degree in hydrology, civil engineering, environmental science, or a related field for the flood risk domain, and a degree in electrical engineering, energy systems, or a related field for the energy infrastructure domain; the resource economics specialist requires a degree in agriculture, agricultural economics, geology, or a related field for the food security domain, and a degree in geology, economic geology, or a related field for the minerals domain. The technical lead requires a university degree in geospatial science, geography, or a related field, with a master's degree preferred for the senior ranks. The office manager requires a university degree and the management qualifications specified by the government's executive rank requirements, with a master's degree in public administration, business administration, or a related field preferred.

The technical certifications are the certifications that establish the officer's competence in the specific tools and methods the office uses, including the geospatial information system (treated in Chapter 29), the remote sensing processing software, and the domain-specific modeling tools. The operators require the certifications for the acquisition desk, the processing engine, the product studio, and the catalog and archive submodules, with the certifications typically offered by the software vendors (the ESRI ArcGIS certifications, the Hexagon Geospatial certifications) or the open-source community (the OSGeo certifications). The specialists require the certifications for the modeling workbench and their domain plug-ins, with the certifications typically offered by the model developers (the HEC certifications for the HEC-HMS and the HEC-RAS models, the crop model certifications for the WOFOST and the DSSAT models, the change detection certifications, the prospectivity certifications). The technical lead requires the certifications for all submodules, with additional certifications in the system configuration and the validation. The office manager requires

the certifications for the workflow management and the commercial release authorization, with the certifications focused on the management rather than the technical operation.

The professional certifications are the certifications offered by the national and international professional bodies, including the certifications in geospatial science, remote sensing, and the domain-specific professions. The geospatial science certifications are offered by the national professional bodies (the American Society for Photogrammetry and Remote Sensing [ASPRS] in the United States, the Surveying and Spatial Sciences Institute [SSSI] in Australia, the Royal Institution of Chartered Surveyors [RICS] in the United Kingdom) and the international professional bodies (the International Society for Photogrammetry and Remote Sensing [ISPRS], the Global Spatial Data Infrastructure Association [GSDI]). The domain-specific certifications are offered by the domain professional bodies (the certified hydrologist certification by the American Institute of Hydrology, the certified crop advisor certification by the American Society of Agronomy, the professional geologist certification by the American Institute of Professional Geologists). The professional certifications are encouraged but not required, and the office supports the officers' pursuit of the professional certifications through the training budget and the study leave.

The qualifications are verified at the appointment and the promotion. The appointment verification checks the officer's academic qualifications (the degree certificates, the transcripts), the technical certifications (the certification records, the expiry dates), and the professional certifications (the certification records, the continuing professional development logs). The promotion verification checks the officer's qualifications for the next rank, with the additional qualifications required for the promotion (the advanced certifications, the leadership certifications, the management qualifications). The verification is conducted by the office manager and the human resources, with the qualifications authority's ratification. The qualifications' verification is the mechanism that ensures the role family's standards, and the verification is the condition of the role family's credibility.

30.4 The Career Paths

The career paths in the geospatial role family are the routes by which an officer progresses through the ranks, and they are designed to produce officers with the breadth and depth the office requires. The three principal career paths are the pipeline path, the production path, and the domain path, with each path providing a different route to the senior ranks and each path producing officers with different strengths.

The pipeline path begins with the data pipeline operator and progresses to the technical lead and the office manager. The data pipeline operator enters the role family at the entry-level operational rank, with the technical qualification in geospatial technology. The operator progresses through the operational ranks, with the promotions based on the demonstrated competence in the data acquisition, the processing, and the base layer maintenance. The operator's career progression may lead to the senior operator rank, and from the senior operator rank, the operator may be promoted to the technical lead, with the bridge rank requiring the additional qualifications in the system configuration and the validation. The technical lead may then be promoted to the office manager, with the executive rank requiring the management qualifications and the competitive selection. The pipeline path produces officers with the deep technical knowledge of the pipeline, and the officers are the office's technical backbone.

The production path begins with the production operator and progresses to the specialist and the technical lead or the office manager. The production operator enters the role family at the entry-level operational rank, with the technical qualification in geospatial technology. The operator progresses through the operational ranks, with the promotions based on the demonstrated competence in the product generation, the classification marking, and the archive management. The operator may then transition to the specialist path, with the transition requiring the additional academic qualification in the domain (the university degree in hydrology, agriculture, geology, or a related field) and the domain-specific certifications. The specialist progresses through the senior operational ranks, with the promotions based on the demonstrated competence in the domain analysis. The specialist may then be promoted to the technical lead (with the bridge rank requiring the system configuration and validation qualifications) or the office manager (with the executive rank requiring the management qualifications and the competitive selection). The production path produces officers with the broad knowledge of the production and the domain, and the officers are the office's analytical backbone.

The domain path begins with the entry through the domain-specific recruitment and progresses to the specialist and the technical lead or the office manager. The domain-specific recruitment recruits the officers with the domain expertise (the hydrology, the agriculture, the geology, or the related field) directly into the specialist rank, with the entry requirements comprising the university degree in the domain and the domain-specific certifications. The specialist progresses through the senior operational ranks, with the promotions based on the demonstrated competence in the domain analysis. The specialist may then be promoted to the technical lead (with the bridge rank requiring the system configuration and validation qualifications) or the office manager (with the executive rank requiring the management qualifications and the competitive selection). The domain path produces officers with the deep domain knowledge, and the officers are the office's analytical specialists.

The three career paths converge at the technical lead and the office manager ranks, and the convergence is the mechanism that allows the office to develop officers with both the technical depth and the institutional breadth. The technical lead may come from any of the three paths, and the technical lead's role requires the technical depth (from the pipeline or the production path) or the domain depth (from the domain path), combined with the system configuration and validation qualifications. The office manager may come from any of the three paths, and the office manager's role requires the technical or the domain depth, combined with the management qualifications and the competitive selection. The convergence is the mechanism that prevents the role family from fragmenting into the three separate careers, and the convergence ensures that the office's leadership has the breadth of experience the leadership requires.

The lateral moves between the paths are encouraged, with the lateral moves providing the officers with the broader experience and the office with the cross-trained officers. The data pipeline operator may move laterally to the production operator, with the move providing the operator with the production experience. The production operator may move laterally to the specialist, with the move requiring the additional academic qualification and the domain-specific certifications. The specialist may move laterally between the hazard and the resource economics specializations, with the move requiring the additional domain-specific certifications and the on-the-job training. The lateral moves are the mechanism that allows the officers to broaden their experience, and the moves are the mechanism that allows the office to develop the cross-trained officers who can serve in multiple roles.

The career paths are the institutional mechanism for the officer's professional development, and the paths are the mechanism by which the office develops its future leadership. The office that maintains clear career paths produces officers who can progress through the ranks and who can eventually lead the office. The office that does not maintain the clear career paths produces officers who stagnate, with the stagnation leading to the attrition and the loss of the institutional knowledge. The career paths are maintained by the office manager, who manages the officers' development, and the qualifications authority, which ratifies the qualifications and the career paths.

30.5 The Language Requirements

The language requirements of the geospatial role family arise from the government's operating language and from the international character of the geospatial profession. The government's operating language is the language in which the government conducts its business, and the geospatial office must produce its products in that language. The international character of the geospatial profession means that the scientific literature, the technical documentation, and the professional discourse are conducted in a small number of international languages (predominantly English, with French, Spanish, and Chinese as the secondary languages), and the geospatial officer must be able to read and understand them. The language requirements are the intersection of these two demands, and they vary by rank and by role.

The operators require the reading and writing proficiency in the government's operating language, because the operators must read the system interfaces (the acquisition desk, the processing engine, the product studio, the catalog and archive) and write the records (the ingestion logs, the processing logs, the production logs) in the operating language. The operators do not require the proficiency in the international languages, because the operators do not read the scientific literature or the technical documentation in the international languages. The operator's language requirement is the basic proficiency (the level B1 or B2 on the Common European Framework of Reference for Languages, or the equivalent on the jurisdiction's language proficiency scale), and the proficiency is verified at the appointment through the language test.

The specialists require the reading and writing proficiency in the operating language, plus the reading proficiency in at least one international language, because the specialists must read the scientific literature and the technical documentation in the international languages. The specialist's reading proficiency in the international language is the advanced proficiency (the level B2 or C1 on the CEFR, or the equivalent), and the proficiency is verified at the appointment through the language test. The specialist's writing proficiency in the international language is not required, because the specialists produce the products in the operating language, not in the international languages. The specialist's international language is typically English, given the dominance of English in the geospatial scientific literature, but the specialist may also meet the requirement through the French, the Spanish, or the Chinese, depending on the jurisdiction's international context.

The technical lead requires the reading and writing proficiency in the operating language, plus the reading and writing proficiency in at least one international language, because the technical lead must read and write the technical documentation in the international languages. The technical lead reads the technical documentation (the software documentation, the model documentation, the standard specifications) and writes the technical documentation (the office's technical standards, the

configuration documentation, the validation reports) in the international language. The technical lead's proficiency in the international language is the advanced proficiency (the level C1 on the CEFR, or the equivalent), and the proficiency is verified at the appointment and the promotion.

The office manager requires the reading, writing, and speaking proficiency in the operating language, plus the reading proficiency in at least one international language, because the manager must brief the desk chief and the bureau chief, must represent the office in the inter-agency coordination, and must read the international literature. The office manager's proficiency in the operating language is the professional proficiency (the level C2 on the CEFR, or the equivalent), and the proficiency includes the speaking proficiency for the briefings, the presentations, and the negotiations. The office manager's proficiency in the international language is the advanced proficiency (the level B2 or C1 on the CEFR, or the equivalent), and the proficiency is verified at the appointment and the promotion.

The language requirements are not static; they evolve as the office's international engagement evolves. The office that engages in the international cooperation (the Copernicus programme, the Group on Earth Observations, the International Charter on Space and Major Disasters) may require the additional language proficiency from its officers, and the office that engages in the bilateral cooperation with the specific countries may require the proficiency in those countries' languages. The language requirements are reviewed periodically by the office manager, with the reviews aligning the requirements with the office's international engagement, and the reviews are ratified by the qualifications authority. The language requirements are the institutional mechanism for the role family's international engagement, and the requirements are the condition of the office's ability to participate in the international geospatial profession.

30.6 The Training and Continuous Development

The training and continuous development of the geospatial role family are the mechanisms by which the officers maintain and update their qualifications, and the mechanisms are the institutional infrastructure for the role family's professional development. The training comprises the initial training (the training provided to the new officers at the appointment), the continuing training (the training provided to the serving officers throughout their careers), and the leadership training (the training provided to the officers preparing for the senior ranks). The continuous development comprises the on-the-job learning, the professional certifications, and the continuing professional development.

The initial training is provided to the new officers at the appointment, with the training comprising the orientation (the office's organization, the pipeline architecture, the role family, the commercial firewall), the technical training (the specific tools and methods the officer will use, including the geospatial information system, the remote sensing processing software, and the domain-specific modeling tools), and the procedural training (the office's procedures, including the ingestion procedures, the processing procedures, the production procedures, the classification marking procedures, and the archive procedures). The initial training is conducted over the first 3 to 6 months of the officer's appointment, with the training delivered by the technical lead, the senior specialists, and the office manager. The initial training's completion is verified through the training assessment, and the assessment's passing is the condition of the officer's confirmation in the role.

The continuing training is provided to the serving officers throughout their careers, with the training comprising the refresher training (the periodic refresher of the existing qualifications), the update training (the training on the new tools, the new methods, and the new procedures), and the advanced training (the training on the advanced topics, the specialized techniques, and the emerging technologies). The continuing training is conducted through the in-house training (the training delivered by the office's senior officers), the external training (the training delivered by the external providers, including the software vendors, the universities, and the professional bodies), and the self-directed learning (the learning conducted by the officers through the online courses, the professional literature, and the conference attendance). The continuing training is supported by the office's training budget, with the budget typically 2 to 5 per cent of the office's personnel budget, and the training is tracked through the officers' training records.

The leadership training is provided to the officers preparing for the senior ranks, with the training comprising the technical leadership training (for the officers preparing for the technical lead rank), the office management training (for the officers preparing for the office manager rank), and the executive training (for the office managers preparing for the higher executive ranks). The technical leadership training covers the system configuration, the validation, the technical standards management, and the team leadership. The office management training covers the office's institutional relationships, the commercial release authorization, the reporting, the budget management, and the staff management. The executive training covers the strategic leadership, the policy engagement, and the inter-agency coordination. The leadership training is typically delivered by the external providers (the government's executive training institute, the university executive programmes, the international leadership programmes), and the training is funded through the office's training budget and the government's executive development budget.

The continuous development comprises the on-the-job learning, the professional certifications, and the continuing professional development. The on-the-job learning is the learning that occurs through the daily work, with the learning facilitated by the technical lead's consultation, the specialists' collaboration, and the office manager's mentoring. The professional certifications are the certifications offered by the professional bodies (treated in section 30.3), with the certifications pursued by the officers and supported by the office's training budget. The continuing professional development is the ongoing learning required by the professional bodies to maintain the certifications, with the development comprising the conference attendance, the publication, the presentation, and the professional service. The continuous development is tracked through the officers' development plans, with the plans reviewed annually by the office manager and the officer.

The training and continuous development are the institutional infrastructure for the role family's professional development, and the infrastructure is the condition of the role family's sustainability. The office that invests in the training and the development produces officers who maintain their qualifications, update their knowledge, and progress through the ranks. The office that does not invest in the training and the development produces officers whose qualifications become outdated, whose knowledge becomes stale, and whose career progression stalls. The training and continuous development are maintained by the office manager (who manages the training budget and the development plans), the technical lead (who manages the technical training), and the qualifications authority (which ratifies the training and the certifications).

30.7 The Payscale Considerations

The payscale considerations for the geospatial role family arise from the labour-market conditions that affect geospatial qualifications. Geospatial skills are in demand in the private sector, in the resource industries, in the technology sector, and in the consulting industry, and the demand means that the government's standard payscale may not be sufficient to attract and retain geospatial officers. The payscale modifier is the mechanism by which the government adjusts the payscale for specific roles where labour-market conditions require it, and the geospatial role family is a candidate for a payscale modifier in most jurisdictions. The modifier is applied as a percentage adjustment to the base payscale, with the percentage calibrated to the labour-market conditions in the jurisdiction.

The payscale modifier's calibration is based on the labour-market analysis, with the analysis comparing the government's payscale to the private-sector payscale for the equivalent qualifications and experience. The analysis is conducted by the office manager (with the human resources' support) and reviewed by the qualifications authority. The analysis considers the private-sector payscale for the geospatial professionals (the surveying firms, the consulting firms, the technology companies, the resource companies), the private-sector payscale for the domain professionals (the hydrology consultants, the agricultural consultants, the geological consultants), and the government's payscale for the equivalent ranks. The analysis identifies the gap between the government's payscale and the private-sector payscale, and the modifier is calibrated to close the gap (typically 10 to 30 per cent, depending on the jurisdiction and the role).

The payscale modifier is applied differently across the role family's seats, with the modifier reflecting the labour-market conditions for each seat. The data pipeline operator and the production operator may receive a moderate modifier (10 to 15 per cent), reflecting the moderate labour-market demand for the geospatial technology professionals. The specialists may receive a higher modifier (15 to 25 per cent), reflecting the higher demand for the domain-specific professionals (the hydrologists, the agricultural scientists, the geologists) and the additional geospatial skills. The technical lead may receive a high modifier (20 to 30 per cent), reflecting the high demand for the senior geospatial professionals with the system configuration and validation expertise. The office manager may receive a moderate modifier (10 to 20 per cent), reflecting the executive payscale's general competitiveness and the office manager's career stage.

The payscale modifier is not a permanent feature of the role family; it is reviewed periodically and adjusted or removed as labour-market conditions change. The review is important because a payscale modifier that is no longer needed is a fiscal cost that the government should not bear, and a payscale modifier that is too low is a recruitment and retention problem that the office should not have to manage. The review is conducted every 2 to 3 years by the office manager (with the human resources' support), with the review comparing the current modifier to the current labour-market conditions and recommending the modifier's adjustment, retention, or removal. The review is ratified by the qualifications authority, and the ratified modifier is applied to the office's payscale.

The payscale modifier is supplemented by the non-salary benefits, with the benefits including the training budget (the office's investment in the officers' professional development, treated in section 30.6), the professional certification support (the office's payment of the certification fees and the study leave for the certification preparation), the conference attendance (the office's funding of the officers' attendance at the professional conferences), and the career progression opportunities (the

office's clear career paths, treated in section 30.4). The non-salary benefits are the mechanism by which the office supplements the payscale and enhances the role family's attractiveness, and the benefits are the condition of the office's ability to attract and retain the officers who may not be motivated solely by the salary.

The payscale modifier and the non-salary benefits are the institutional mechanisms for the geospatial role family's labour-market competitiveness, and the mechanisms are the condition of the office's ability to attract and retain the officers it needs. An office that maintains the competitive payscale and the attractive benefits builds a role family that attracts and retains the qualified officers; an office that does not struggles to attract and retain the officers, with the struggles leading to the vacancies, the workloads, and the attrition. The payscale considerations are the institutional framework for the role family's labour-market position, and the framework is the condition of the role family's viability.

30.8 The Comparative Practice

The comparative practice of geospatial role families is examined through the national civil service classifications for geospatial professionals, the professional certification frameworks, and the payscale modifiers. The comparative practice spans three jurisdictions — the United States, Australia, and France — with the three jurisdictions representing three variants of the institutional arrangement and the qualifications framework.

The United States model is exemplified by the federal government's classification of geospatial professionals in the General Schedule (GS). The geospatial professionals are classified in the GS-0300 series (the cartography series, GS-1370; the geography series, GS-0150; the geology series, GS-1350), with the United States Geological Survey (USGS) and the National Geospatial-Intelligence Agency (NGA) as the principal employers. The USGS employs approximately 8,000 staff, with the geospatial professionals distributed across the GS-0300 series and the related series (the hydrology series, GS-1315; the ecology series, GS-0408). The NGA employs approximately 14,500 staff, with the geospatial professionals distributed across the GS-0300 series and the intelligence series. The qualifications for the GS-0300 series comprise the academic qualifications (the bachelor's degree in the relevant field, with the master's degree preferred for the senior ranks), the technical certifications (the vendor certifications, the professional certifications), and the security clearance (the Secret or Top Secret clearance for the NGA positions). The payscale is the GS payscale, with the locality pay adjustment (the cost-of-living adjustment for the high-cost areas) and the special pay rates for the hard-to-fill positions (the geospatial positions in the NGA may receive the special pay rates).

The Australian model is exemplified by the federal government's classification of geospatial professionals in the spatial information series. The geospatial professionals are classified in the Australian Public Service (APS) classification, with the spatial information roles at the APS-1 to APS-6 levels (the operational tier) and the EL-1 to EL-2 levels (the executive tier). The Geoscience Australia (GA) and the state land departments (such as the New South Wales Land and Property Information, the Victorian Department of Environment, Land, Water and Planning) are the principal employers. The GA employs approximately 600 staff, with the geospatial professionals distributed across the spatial information roles and the related roles (the geology roles, the hydrology roles). The

qualifications comprise the academic qualifications (the bachelor's degree in the relevant field), the technical certifications (the vendor certifications, the professional certifications), and the professional certifications (the SSSI certifications). The payscale is the APS payscale, with the market loading for the hard-to-fill positions (the geospatial positions may receive the market loading of 10 to 20 per cent).

The French model is exemplified by the civil service's classification of geospatial professionals in the engineers and technicians corps. The geospatial professionals are classified in the corps of the engineers (the Ingénieurs de l'État) and the technicians (the Techniciens de l'État), with the Institut national de l'information géographique et forestière (IGN, the national geographic institute) as the principal employer. The IGN employs approximately 1,800 staff, with the geospatial professionals distributed across the engineers corps (the Ingénieurs des travaux géographiques et forestiers, the Ingénieurs de recherche) and the technicians corps (the Techniciens supérieurs de l'information géographique et forestière). The qualifications comprise the academic qualifications (the engineering degree from the Grandes Écoles, the master's degree from the universities), the technical certifications (the internal certifications), and the professional certifications (the Ordre des géomètres-experts certifications). The payscale is the civil service payscale, with the prime de fonctions et de résultats (the performance bonus) and the indemnité de résidence (the residence allowance) for the specific positions.

The three models share the role family's structure: the operational and executive tiers, the domain specializations, the academic and technical and professional qualifications, the career paths, the language requirements, and the payscale modifiers. The three models also share the payscale challenge: the geospatial qualifications command a labour-market premium in most jurisdictions, and the governments that do not adjust their payscales struggle to attract and retain geospatial professionals. The three models' payscale adjustments (the U.S. special pay rates, the Australian market loading, the French performance bonus) are the mechanisms by which the governments address the challenge, and the mechanisms are the variants of the payscale modifier treated in section 30.7.

The comparative practice also includes the professional certification frameworks. The ASPRS (the American Society for Photogrammetry and Remote Sensing) offers the Certified Photogrammetrist, the Certified Mapping Scientist, and the Certified GIS/LIS Technologist certifications in the United States. The SSSI (the Surveying and Spatial Sciences Institute) offers the Certified Practising Surveyor, the Certified Hydrographic Surveyor, and the Certified GIS Professional certifications in Australia. The RICS (the Royal Institution of Chartered Surveyors) offers the Chartered Surveyor certifications in the United Kingdom. The ISPRS (the International Society for Photogrammetry and Remote Sensing) offers the international certifications. The geospatial role family's professional certifications are aligned with these frameworks, and the officers' pursuit of the certifications is supported by the office's training budget and the study leave. The professional certifications are the international recognition of the geospatial profession, and the certifications are the condition of the role family's international mobility.

30.9 Worked Example: The Recruitment and Career Progression of a Hazard Specialist

The recruitment and career progression is illustrated through the example of a hazard specialist who joins the office in 2018 and progresses through the ranks to the technical lead by 2026. The example demonstrates the role family's operation across the appointment, the qualifications, the career path, the training, and the promotion, and the example is the illustration of the role family's institutional framework.

The recruitment begins in March 2018, when the office advertises the hazard specialist position following the retirement of the previous senior hazard specialist. The advertisement specifies the academic qualifications (the bachelor's or the master's degree in hydrology, civil engineering, environmental science, or a related field), the technical certifications (the HEC certifications for the HEC-HMS and the HEC-RAS models, or the equivalent), the professional certifications (the Certified Hydrologist certification by the American Institute of Hydrology, or the equivalent, preferred but not required), the language requirements (the reading and writing proficiency in the operating language, plus the reading proficiency in at least one international language), and the experience (3 to 5 years of the hydrological modelling experience, preferably in the government or the consulting sector). The advertisement is published on the government's recruitment portal and the professional bodies' job boards (the ASPRS job board, the SSSI job board, the American Institute of Hydrology job board).

The application period closes in April 2018, with 27 applications received. The applications are reviewed by the office manager and the technical lead, with the review assessing the academic qualifications, the technical certifications, the professional experience, and the language proficiency. The review shortlists 6 candidates for the interview, with the shortlisting based on the qualifications' match, the experience's relevance, and the certifications' currency. The interviews are conducted in May 2018, with the interview panel comprising the office manager, the technical lead, and the senior resource economics specialist (as the cross-domain representative). The interview assesses the candidate's technical knowledge (the hydrological modelling, the hydraulic modelling, the flood risk assessment), the analytical judgment (the interpretation of the model results, the validation against the observed data), the communication skills (the product preparation, the consumer briefing), and the institutional fit (the understanding of the government's geospatial function, the commercial firewall, the role family).

The successful candidate, Dr. Maren Holloway, is offered the position in June 2018 and accepts the offer. Dr. Holloway holds a Ph.D. in hydrology from the University of Portsmith (2014), with the dissertation on the flood hydrology of the Riverbourne Basin. She holds the HEC certifications (the HEC-HMS and the HEC-RAS certifications, obtained in 2015), the Certified Hydrologist certification (obtained in 2017), and the reading proficiency in English (the international language) at the C1 level. She has 4 years of the hydrological modelling experience, with the previous employment at the Riverbourne Water Authority (2014-2018) as the hydrologist. Dr. Holloway's qualifications and experience meet the position's requirements, and her interview performance demonstrates the technical knowledge, the analytical judgment, the communication skills, and the institutional fit.

The initial training is conducted in July-August 2018, with the training comprising the orientation (the office's organization, the pipeline architecture, the role family, the commercial firewall, delivered by the office manager), the technical training (the geospatial information system, the modeling workbench, the hydrology plug-in, delivered by the technical lead), and the procedural training (the ingestion procedures, the processing procedures, the production procedures, the classification marking procedures, the archive procedures, delivered by the senior officers). The initial training's completion is verified through the training assessment in September 2018, with the assessment comprising the written test (the technical knowledge and the procedures), the practical test (the model run and the product generation), and the oral test (the institutional understanding and the communication). Dr. Holloway passes the assessment and is confirmed in the hazard specialist role at the entry-level operational rank (the rank equivalent to the APS-5 or the GS-11).

The continuing training is conducted throughout Dr. Holloway's career. In 2019, she attends the HEC-RAS 2D advanced training (the 5-day course delivered by the Hydrologic Engineering Center), with the training funded by the office's training budget. In 2020, she attends the European Flood Awareness System (EFAS) training (the 3-day course delivered by the European Commission's Joint Research Centre), with the training funded by the office's training budget and the Copernicus programme's fellowship. In 2021, she completes the ASPRS Certified Mapping Scientist certification (the certification requiring the experience, the examination, and the continuing professional development), with the certification fees and the study leave funded by the office's training budget. In 2022, she attends the International Association of Hydrological Sciences (IAHS) conference in Portsmith, with the conference attendance funded by the office's training budget, and she presents the paper on the Riverbourne Basin flood modelling (the paper co-authored with the technical lead).

The career progression begins with the promotion to the senior operational rank in 2022, after 4 years of service. The promotion is based on the demonstrated competence in the flood risk and the energy infrastructure domains, the advanced certifications (the HEC-RAS 2D, the ASPRS CMS), the conference presentation, and the office manager's recommendation. The promotion is ratified by the qualifications authority, and Dr. Holloway's payscale is adjusted to the senior operational rank (the rank equivalent to the APS-6 or the GS-13), with the payscale modifier of 18 per cent (the modifier for the senior hazard specialist, calibrated to the labour-market conditions).

The career progression continues with the promotion to the technical lead in 2026, after 8 years of service. The promotion is based on the demonstrated leadership in the technical standards (the configuration management, the validation, the technical standards review), the additional qualifications (the system configuration and validation certifications, obtained in 2024-2025), the leadership training (the technical leadership training, completed in 2025), and the competitive selection (the selection through the office's competitive process, with the panel comprising the office manager, the bureau chief, and the qualifications authority's representative). The promotion is ratified by the qualifications authority, and Dr. Holloway's payscale is adjusted to the bridge rank (the rank equivalent to the EL-1 or the GS-14), with the payscale modifier of 25 per cent (the modifier for the technical lead, calibrated to the labour-market conditions).

The career progression's lateral moves include the temporary assignment to the resource economics specialization in 2023, with the assignment providing Dr. Holloway with the cross-domain experience. The assignment lasts 3 months (the March-May 2023 period), during which Dr.

Holloway works alongside the resource economics specialist on the Northfields drought assessment (the assessment treated in Chapter 25, section 25.9). The assignment provides Dr. Holloway with the exposure to the crop model plug-in, the supply-demand analysis, and the procurement advisory production, and the exposure broadens her analytical experience and prepares her for the technical lead role's cross-domain responsibilities.

The career progression's professional development includes the continuing professional development for the Certified Hydrologist and the ASPRS CMS certifications. The CPD comprises the 40 hours of the continuing education per year (the conference attendance, the workshop participation, the online courses), the publication (the paper in the Journal of Hydrology in 2024, the paper in the Photogrammetric Engineering and Remote Sensing in 2025), and the professional service (the review for the Journal of Hydrology, the committee membership in the IAHS). The CPD is tracked through Dr. Holloway's development plan, with the plan reviewed annually by the office manager.

This worked example illustrates the role family's full operation: the recruitment (the advertisement, the application, the interview, the offer, the acceptance), the initial training (the orientation, the technical training, the procedural training, the assessment, the confirmation), the continuing training (the HEC-RAS 2D training, the EFAS training, the ASPRS CMS certification, the IAHS conference), the career progression (the promotion to the senior operational rank in 2022, the promotion to the technical lead in 2026), the lateral move (the temporary assignment to the resource economics specialization in 2023), the professional development (the CPD, the publications, the professional service), and the payscale adjustments (the 18 per cent modifier for the senior specialist, the 25 per cent modifier for the technical lead). The role family's institutional framework operated as designed: the recruitment attracted the qualified candidate, the initial training prepared her for the role, the continuing training maintained and updated her qualifications, the career progression provided the advancement opportunities, the lateral move broadened her experience, and the payscale adjustments maintained the labour-market competitiveness. The framework is the condition of the office's ability to attract, develop, and retain the officers it needs, and the framework is the reason the role family exists. Dr. Holloway's career, from the recruitment in 2018 to the technical lead promotion in 2026, is the illustration of the role family's operation, and the career is the example that this chapter provides for the qualifications authority's ratification. The role family that is ratified becomes a permanent feature of the government's rank structure, and the ratification is the condition of the role family's permanence, and the permanence is the condition of the office's ability to serve the jurisdiction's geospatial intelligence needs over the long term.

Chapter 31 — Training and the Administrative Training Authority

31.1 The Training Pipeline

Training is delivered through the administrative training authority's course catalogue, comprising the foundational course, the submodule courses, the domain courses, and the leadership courses, forming the training pipeline from the new hire to the certified specialist and the senior leader. The training pipeline is rank-progressive and role-specific, following the model of the world's leading government training institutions, and it is the mechanism by which the office develops the officers' competence in the geospatial profession. The pipeline is the institutional expression of the office's commitment to the officers' professional development, and the pipeline is the condition of the office's ability to produce the intelligence the office's consumers require.

The training pipeline is rank-progressive, meaning that the courses are sequenced to match the officer's progression through the ranks, with the foundational courses for the entry-level officers, the submodule and the domain courses for the operational officers, the advanced courses for the senior officers, and the leadership courses for the officers preparing for the executive ranks. The pipeline is role-specific, meaning that the courses are tailored to the officer's role, with the pipeline courses for the data pipeline operator, the production courses for the production operator, the hazard domain courses for the hazard specialist, and the resource economics domain courses for the resource economics specialist. The rank-progressive and role-specific design ensures that the training is relevant to the officer's work and that the training builds the officer's competence over the career.

The training pipeline's architecture has four stages, each corresponding to a phase of the officer's career. The first stage is the foundational stage, comprising the foundational course that all new officers complete within the first 6 months of the appointment. The second stage is the operational stage, comprising the submodule and the domain courses that the officers complete within the first 12 to 18 months of the appointment, leading to the operational certification. The third stage is the advanced stage, comprising the advanced courses that the senior officers complete within the 5 to 8 years of the appointment, leading to the advanced certification. The fourth stage is the leadership stage, comprising the leadership courses that the officers complete before the promotion to the technical lead or the office manager, leading to the leadership certification. The four stages together constitute the training pipeline, and the pipeline is the institutional infrastructure for the officer's professional development.

The training pipeline is delivered by the administrative training authority, which is the government's training institution that serves the whole government. The authority maintains the course catalogue, develops the courses, delivers the courses, and awards the certifications. The geospatial office partners with the authority for the training, with the office providing the domain-specific content (the hydrology, the crop modelling, the change detection, the prospectivity) and the authority providing the training infrastructure (the course management, the instructors, the assessment, the certification). The partnership is the relationship between a client and a provider, with the office as the client and the authority as the provider, and the partnership is the mechanism by which the geospatial training pipeline is delivered.

31.2 The Foundational Course

The foundational course is the entry-level course that all new geospatial officers complete, regardless of their role. The course comprises five modules: the introduction to the geospatial profession, the introduction to the pipeline architecture, the introduction to the information system, the introduction to the classification framework, and the introduction to the office's procedures. The course is the foundation on which the subsequent role-specific courses build, and the course is the mechanism by which the new officer acquires the common knowledge that all geospatial officers share. The course is delivered by the administrative training authority, with the office's technical lead and the office manager contributing the office-specific content, and the course is the prerequisite for the submodule and the domain courses.

The introduction to the geospatial profession module covers the profession's history, the profession's institutions, and the profession's standards. The history covers the evolution of the geospatial intelligence from the military reconnaissance to the civilian earth observation, the establishment of the national mapping agencies and the geological surveys in the 19th century, the Landsat programme's launch in 1972, and the Copernicus programme's establishment in 2014. The institutions cover the national and the international geospatial institutions (the USGS, the ESA, the JRC, the GEO, the CEOS), the professional bodies (the ASPRS, the ISPRS, the SSSI, the RICS), and the standards organizations (the OGC, the ISO TC 211). The standards cover the technical standards (the OGC standards for the web services, the ISO 19100 series for the geographic information), the professional standards (the ASPRS accuracy standards, the ISPRS commission standards), and the ethical standards (the professional ethics, the confidentiality, the conflict of interest).

The introduction to the pipeline architecture module covers the pipeline's six capabilities (the base layer, the acquisition, the processing, the modeling, the production, the archive, treated in Part II), the pipeline's shared-backbone and swappable-heads design (treated in Chapter 6 and Chapter 29), and the pipeline's four domains (the flood risk, the food security, the energy infrastructure, the minerals, treated in Part IV). The module is the conceptual foundation for the officer's understanding of the office's work, and the module is the prerequisite for the submodule courses that cover the pipeline's capabilities in detail.

The introduction to the information system module covers the information system's six submodules (the spatial data infrastructure, the acquisition desk, the processing engine, the modeling workbench, the product studio, the catalog and archive, treated in Chapter 29), the information system's integration with the government's core administrative information system (the records management, the financial management, the human resources, the procurement, the workflow management), and the information system's technology architecture (the geospatial database, the web services, the computing infrastructure). The module is the technical foundation for the officer's understanding of the office's tools, and the module is the prerequisite for the submodule courses that cover the submodules' operation in detail.

The introduction to the classification framework module covers the classification's four levels (the Public, the Internal, the Confidential, the Security Confidential, treated in Chapter 11, section 11.3), the classification's application to the office's products (the classification marking, the access control, the distribution), the classification's legal framework (the official secrets, the critical infrastructure information protection, the commercial confidentiality), and the classification's

enforcement (the production operator's marking, the office manager's review, the audit). The module is the security foundation for the officer's understanding of the office's classification discipline, and the module is the prerequisite for the officer's handling of the classified products.

The introduction to the office's procedures module covers the office's daily procedures (the morning ingestion check, the processing queue review, the production run, the classification marking verification, the archive update, treated in Part III), the office's event-driven procedures (the activation, the event-driven production, the stand-down, treated in Chapter 13), the office's commercial procedures (the commercial charge code management, the commercial release authorization, treated in Chapter 5), and the office's reporting procedures (the monthly reporting, the annual reporting, treated in Chapter 13). The module is the operational foundation for the officer's understanding of the office's daily work, and the module is the prerequisite for the officer's participation in the office's operations.

31.3 The Submodule Courses

The submodule courses are the role-specific courses that train the officers in the operation of the information system's submodules. The courses comprise the spatial data infrastructure course, the acquisition desk course, the processing engine course, the modeling workbench course, the product studio course, and the catalog and archive course. Each course covers the submodule's operation, configuration, and maintenance, and each course is tailored to the role that primarily operates the submodule. The data pipeline operator takes the spatial data infrastructure, the acquisition desk, and the processing engine courses; the production operator takes the product studio and the catalog and archive courses; the technical lead takes all six submodule courses, with the additional configuration and validation content; and the specialists take the modeling workbench course, with the domain-specific content delivered through the domain courses.

The spatial data infrastructure course covers the base layer storage (the PostGIS database, the spatial indexing, the spatial functions, the spatial data types, treated in Chapter 29, section 29.2), the base layer versioning (the valid-from and valid-to timestamps, the temporal queries, the AS OF and BETWEEN queries), the base layer synchronization (the incremental updates, the conflict resolution, the monitoring), and the base layer access interface (the OGC WMS, WFS, WCS, the database connections, the monitoring). The course is delivered over 5 days, with the morning sessions covering the theory and the afternoon sessions covering the practical exercises. The course's assessment comprises the written test (the theory) and the practical test (the base layer operation, including the ingestion, the synchronization, and the access interface's configuration). The course is the prerequisite for the data pipeline operator's certification in the spatial data infrastructure operation.

The acquisition desk course covers the tasking request interface (the satellite tasking APIs, the tasking request submission, the status tracking, treated in Chapter 29, section 29.3), the purchase order interface (the budget validation, the procurement rules, the license validation), the sensor feed subscription interface (the Sentinel-1, 2, 3, the Landsat 8, 9, the subscription management), the license record management (the license types, the license terms, the license enforcement), and the ingestion verification (the completeness, the format, the georeferencing, the quality). The course is delivered over 5 days, with the morning sessions covering the theory and the afternoon sessions

covering the practical exercises. The course's assessment comprises the written test and the practical test (the acquisition desk operation, including the tasking request, the purchase order, the subscription management, and the license record management).

The processing engine course covers the processing queue management (the workflow engine, the job prioritization, the resource allocation, the status tracking), the processing parameter management (the georeferencing, the calibration, the mosaicking, the projection, the cloud masking, the data fusion parameters, treated in Chapter 9), the ground control point library (the points' coordinates, accuracy, description, source), the processing output catalog (the product's identifier, type, source, parameters, quality, storage), and the validation interface (the geometric accuracy, the radiometric accuracy, the completeness). The course is delivered over 5 days, with the assessment comprising the written test and the practical test (the processing engine operation, including the queue management, the parameter configuration, the validation).

The modeling workbench course covers the workbench environment (the data access interface, the computing resources, the model execution engine, the model output interface, treated in Chapter 29, section 29.5), the domain plug-ins (the hydrology plug-in, the crop model plug-in, the change detection plug-in, the prospectivity plug-in, treated in Chapter 10), the parameter management (the model parameters, the parameter library, the versioning), the validation framework (the validation tools, the validation against the observed data), and the provenance recording (the model version, the parameters, the run date, the validation results, the output storage). The course is delivered over 5 days, with the assessment comprising the written test and the practical test (the modeling workbench operation, including the plug-in execution, the parameter configuration, the validation, the provenance recording). The specialists take the modeling workbench course in conjunction with their domain courses, with the workbench course providing the generic workbench operation and the domain courses providing the domain-specific plug-in operation.

The product studio course covers the cartographic template management (the templates, the layout, the scale, the orientation, the size, treated in Chapter 29, section 29.6), the symbology library (the symbology by domain, the symbology by feature type), the classification marking tools (the markings, the format, the locations, the verification), and the format conversion tools (the GeoTIFF, the Shapefile, the GeoPackage, the KML, the PDF, the DOCX, the format-specific parameters, the conversion verification). The course is delivered over 5 days, with the assessment comprising the written test and the practical test (the product studio operation, including the template application, the symbology, the classification marking, the format conversion).

The catalog and archive course covers the product catalog (the catalog records, the UUIDs, the access permissions, treated in Chapter 29, section 29.7), the version history (the version creation, the version retrieval, the auditability), the classification markings (the markings, the access control, the enforcement), the retention designations (the retention periods, the dispositions, the records management), the access logging (the logging, the review, the compliance), and the provenance chain (the chain's links, the chain's verification, the auditability). The course is delivered over 5 days, with the assessment comprising the written test and the practical test (the catalog and archive operation, including the cataloging, the version management, the access control, the retention management, the provenance chain).

31.4 The Domain Courses

The domain courses are the specialist courses that train the specialists in the domain-specific analytical methods. The courses comprise the hydrology course for the flood risk domain, the crop model course for the food security domain, the change detection course for the energy infrastructure domain, and the prospectivity course for the minerals domain. Each course covers the domain's methods, models, validation protocols, and product types, and each course is tailored to the specialist who operates the domain's modeling plug-in. The hazard specialist takes the hydrology and the change detection courses; the resource economics specialist takes the crop model and the prospectivity courses.

The hydrology course covers the rainfall-runoff modelling (the HEC-HMS model, the GR4J model, the curve number method, the unit hydrograph method, treated in Chapter 10, section 10.2), the hydraulic modelling (the HEC-RAS 2D model, the Delft3D model, the 1D and 2D modelling, the boundary conditions), the inundation mapping (the post-processing of the hydraulic model's outputs, the depth grids, the arrival times), the flood frequency analysis (the annual exceedance probability, the return periods, the Gumbel and the GEV distributions), and the validation protocols (the comparison with the observed stages, the comparison with the observed extents, the comparison with the observed durations). The course is delivered over 10 days (the 5-day HEC-HMS module and the 5-day HEC-RAS module), with the assessment comprising the written test and the practical test (the hydrology modelling, including the HEC-HMS run, the HEC-RAS run, the inundation mapping, the validation).

The crop model course covers the vegetation index computation (the NDVI, the EVI, the NDWI, the computation from the Sentinel-2 surface reflectance, treated in Chapter 10, section 10.3), the phenological stage determination (the crop calendar, the accumulated growing degree days, the stage transitions), the yield forecasting (the statistical regression, the mechanistic models with the WOFOST and the DSSAT, the calibration, the validation), the supply-demand analysis (the production estimate, the consumption estimate, the balance, the release or buy signal), and the validation protocols (the comparison with the reported yields, the comparison with the observed conditions, the comparison with the observed stages). The course is delivered over 10 days (the 5-day vegetation and phenology module and the 5-day yield forecasting and supply-demand module), with the assessment comprising the written test and the practical test (the crop modelling, including the vegetation index computation, the phenological stage determination, the yield forecasting, the supply-demand analysis, the validation).

The change detection course covers the change detection approaches (the image differencing, the object-based change detection, the threshold selection, the segmentation, treated in Chapter 10, section 10.4), the encroachment classification (the vegetation encroachment, the third-party construction, the ground subsidence, the excavation), the vulnerability scoring (the criticality weight, the proximity factor, the change type factor, the score computation), the high-resolution imagery analysis (the Maxar WorldView-3, the Airbus Pleiades Neo, the tasking, the processing), and the validation protocols (the comparison with the field-verified changes, the comparison with the actual failure occurrences, the comparison with the field classifications). The course is delivered over 10 days (the 5-day change detection module and the 5-day vulnerability scoring and validation module), with the assessment comprising the written test and the practical test (the change detection, including

the image differencing, the object-based detection, the encroachment classification, the vulnerability scoring, the validation).

The prospectivity course covers the knowledge-driven prospectivity (the boolean logic, the fuzzy logic, the weights of evidence, treated in Chapter 10, section 10.5), the data-driven prospectivity (the logistic regression, the random forest, the neural networks), the resource estimation (the kriging, the inverse distance weighting, the resource classification), the geological data analysis (the geological maps, the geochemical data, the geophysical data), and the validation protocols (the comparison with the actual mineral occurrences, the comparison with the drilling results, the comparison with the prediction confidence). The course is delivered over 10 days (the 5-day prospectivity module and the 5-day resource estimation and validation module), with the assessment comprising the written test and the practical test (the prospectivity modelling, including the knowledge-driven and the data-driven prospectivity, the resource estimation, the validation).

The domain courses are the courses that build the specialists' analytical competence, and the courses are the mechanism by which the specialists learn to produce the domain's intelligence. The courses are delivered by the administrative training authority, with the office's technical lead providing the workbench-specific configuration and the office manager providing the office-specific product standards. The courses are the prerequisite for the specialists' certification in the domain operation, and the certification is the qualification that the specialist is competent to operate the domain's modeling plug-in. The domain courses are the analytical foundation of the specialists' competence, and the courses are the condition of the office's intelligence quality.

31.5 The Leadership Courses

The leadership courses are the courses that train the officers preparing for the senior ranks, with the courses comprising the technical leadership course, the office management course, and the executive course. The courses are the fourth stage of the training pipeline, and the courses are the prerequisite for the promotion to the technical lead and the office manager. The courses are delivered by the administrative training authority, with the office's office manager and the bureau chief contributing the office-specific content.

The technical leadership course covers the system configuration (the geospatial information system's configuration, the parameter management, the version control, treated in Chapter 29), the validation (the validation framework, the validation against the independent data, the validation reporting), the technical standards management (the standards' development, the standards' review, the standards' documentation), the team leadership (the team's organization, the team's mentoring, the team's performance management), and the technical project management (the project's planning, the project's execution, the project's reporting). The course is delivered over 10 days (the 5-day technical module and the 5-day leadership module), with the assessment comprising the written test, the practical test (the system configuration and the validation), and the oral test (the team leadership and the project management). The course is the prerequisite for the technical lead certification, and the certification is the qualification for the promotion to the technical lead rank.

The office management course covers the office's institutional relationships (the desk chief, the bureau chief, the inter-agency partners, treated in Chapter 13), the commercial release authorization (the commercial firewall, the release procedure, the four disciplines' interaction, treated in Chapter 5),

the reporting (the monthly reporting, the annual reporting, the audit), the budget management (the appropriation, the commercial revenue, the cost accounting, the training budget), the staff management (the recruitment, the development, the performance review, the discipline), and the inter-agency coordination (the emergency management authority, the strategic reserve authority, the sectoral regulator, the water authority, the national statistics agency, treated in Part V). The course is delivered over 10 days, with the assessment comprising the written test, the practical test (the commercial release authorization and the reporting), and the oral test (the institutional relationships and the inter-agency coordination). The course is the prerequisite for the office manager certification, and the certification is the qualification for the promotion to the office manager rank.

The executive course covers the strategic leadership (the strategic planning, the strategic execution, the strategic review), the policy engagement (the policy development, the policy advice, the policy advocacy's prohibition), the inter-agency coordination at the executive level (the executive committees, the inter-ministerial coordination, the cabinet briefings), and the government's executive framework (the cabinet system, the budget cycle, the accountability framework). The course is delivered over 5 to 10 days, with the assessment comprising the written test and the oral test (the strategic leadership and the policy engagement). The course is the prerequisite for the executive certification, and the certification is the qualification for the higher executive ranks.

The leadership courses are the courses that build the officers' leadership competence, and the courses are the mechanism by which the office develops its future leadership. The courses are delivered by the administrative training authority, with the office's office manager and the bureau chief contributing the office-specific content, and the courses are often delivered in partnership with the government's executive training institute or the university executive programmes. The leadership courses are the leadership foundation of the senior officers' competence, and the courses are the condition of the office's leadership succession.

31.6 The Continuing Education

The continuing education is the ongoing training that the officers complete throughout their careers, and it is the mechanism by which the officers maintain and update their competence. The continuing education comprises the refresher courses, the update courses, and the advanced courses, with each course type serving a different purpose in the officer's professional development.

The refresher courses refresh the officers' competence in the foundational, the submodule, and the domain courses. The refresher courses are delivered every 2 to 3 years, with the courses covering the same content as the original courses but with the updated examples and the updated exercises. The refresher courses are the mechanism by which the officers maintain their original certifications, and the completion of the refresher courses is the condition of the certifications' currency. The refresher courses are typically delivered over 2 to 3 days (shorter than the original 5-day courses), with the assessment comprising the written test and the practical test.

The update courses update the officers' competence in the new methods, the new tools, and the new procedures. The update courses are delivered as the new methods, the new tools, and the new procedures are introduced, with the courses covering the specific updates and the migration from the old to the new. The update courses are the mechanism by which the officers adopt the new methods, the new tools, and the new procedures, and the completion of the update courses is the condition of

the officers' currency. The update courses are typically delivered over 1 to 3 days, with the assessment comprising the practical test (the new method's application, the new tool's operation, the new procedure's execution).

The advanced courses advance the officers' competence in the specialized topics, the emerging methods, and the leadership skills. The advanced courses are delivered to the senior officers, with the courses covering the specialized topics (the ensemble forecasting, the machine learning, the cloud computing), the emerging methods (the InSAR, the UAV photogrammetry, the deep learning), and the leadership skills (the project management, the team leadership, the change management). The advanced courses are the mechanism by which the officers deepen their competence and prepare for the senior ranks, and the completion of the advanced courses is the condition of the promotion to the senior ranks. The advanced courses are typically delivered over 5 to 10 days, with the assessment comprising the written test, the practical test, and the oral test.

The continuing education is not optional; it is a requirement of the officer's certification, and the certification must be maintained through the continuing education. The officer who does not complete the continuing education loses the certification, and the loss of the certification means the officer is no longer qualified to operate the submodule or the domain. The continuing education is tracked through the officers' training records, with the records specifying the courses completed, the certifications maintained, and the continuing education hours accumulated. The training records are reviewed annually by the office manager, with the review identifying the officers who are due for the refresher or the update courses and the officers who are eligible for the advanced courses. The continuing education is maintained by the office manager (who tracks the officers' continuing education) and the administrative training authority (which delivers the courses).

The continuing education is supplemented by the professional development activities, with the activities comprising the conference attendance, the publication, the presentation, and the professional service. The conference attendance is the officers' attendance at the professional conferences (the ISPRS Congress, the ESA Living Planet Symposium, the AGU Fall Meeting, the FIG Working Week), with the attendance funded by the office's training budget. The publication is the officers' publication of the research papers in the professional journals (the Photogrammetric Engineering and Remote Sensing, the ISPRS Journal of Photogrammetry and Remote Sensing, the International Journal of Applied Earth Observation and Geoinformation), with the publication supported by the office's research time. The presentation is the officers' presentation at the professional conferences and the workshops, with the presentation prepared during the office's research time. The professional service is the officers' service on the professional bodies' committees and the editorial boards, with the service recognized as the professional contribution. The professional development activities are tracked through the officers' development plans, with the plans reviewed annually by the office manager.

31.7 The Certification

The certification is the qualification that the officer is competent to operate a submodule or a domain, and the certification is the institutional mechanism for the officer's professional standing. The certification is awarded by the administrative training authority, on the completion of the course and the passing of the assessment, and the certification is maintained through the continuing

education. The certification is the qualification that the office manager and the qualifications authority rely on, and the certification is the condition of the officer's assignment to the role. An officer who is not certified is not qualified, and an officer who is not qualified must not operate the submodule or the domain.

The certification's structure comprises the foundational certification, the submodule certifications, the domain certifications, and the leadership certifications. The foundational certification is awarded on the completion of the foundational course and the passing of the foundational assessment, and the certification is the qualification for the officer's confirmation in the role. The submodule certifications are awarded on the completion of the submodule courses and the passing of the submodule assessments, with the certifications specifying the submodule (the spatial data infrastructure certification, the acquisition desk certification, the processing engine certification, the modeling workbench certification, the product studio certification, the catalog and archive certification). The domain certifications are awarded on the completion of the domain courses and the passing of the domain assessments, with the certifications specifying the domain (the hydrology certification, the crop model certification, the change detection certification, the prospectivity certification). The leadership certifications are awarded on the completion of the leadership courses and the passing of the leadership assessments, with the certifications specifying the leadership level (the technical leadership certification, the office management certification, the executive certification).

The certification's assessment is rigorous, with the assessment comprising the written test, the practical test, and (for the leadership certifications) the oral test. The written test covers the theory, with the multiple-choice questions, the short-answer questions, and the essay questions. The practical test covers the operation, with the officer's execution of the submodule's or the domain's operation under the assessor's observation. The oral test covers the leadership, with the officer's response to the leadership scenarios under the panel's questioning. The assessment's passing threshold is typically 75 per cent, with the candidates who score below the threshold required to retake the course and the assessment. The assessment's rigour is the condition of the certification's credibility, and the rigour is the mechanism by which the certification maintains its standards.

The certification's maintenance is the ongoing requirement, with the certification maintained through the continuing education. The certification is valid for 3 years, with the recertification requiring the completion of the refresher courses and the accumulation of the continuing education hours (typically 40 hours per year, or 120 hours over the 3-year recertification cycle). The certification that is not maintained through the continuing education expires, and the expiry means the officer is no longer qualified to operate the submodule or the domain. The certification's maintenance is tracked through the officers' training records, with the records specifying the certification's award date, the certification's expiry date, and the continuing education hours accumulated. The maintenance is verified by the office manager (who tracks the officers' certifications) and the qualifications authority (which ratifies the certifications and the recertifications).

The certification's recognition is the institutional acceptance of the certification by the other government agencies, the professional bodies, and the international organizations. The certification is recognized by the other government agencies through the government's qualifications framework,

with the certification's equivalence to the other agencies' certifications established through the framework. The certification is recognized by the professional bodies through the articulation agreements, with the certification's equivalence to the professional certifications (the ASPRS, the SSSI, the RICS) established through the agreements. The certification is recognized by the international organizations through the mutual recognition arrangements, with the certification's equivalence to the international certifications (the ISPRS, the OGC) established through the arrangements. The certification's recognition is the condition of the officers' mobility across the government, the profession, and the international community, and the recognition is the mechanism by which the certification's value extends beyond the office.

31.8 The Comparative Practice

The comparative practice of the government training institutions is examined through five models, with the five models representing the variants of the institutional arrangement and the training pipeline's design. The five models are the United States Foreign Service Institute, the Canada School of Public Service, the French École nationale d'administration, the German Bundesakademie für öffentliche Verwaltung, and the Singapore Civil Service College.

The United States Foreign Service Institute (FSI) is the training institution for the United States foreign service, established in 1947 and located at the George P. Shultz National Foreign Affairs Training Center in Arlington, Virginia. The FSI offers the foundational courses for the entry-level officers (the A-100 course for the foreign service officers, the specialist courses for the foreign service specialists), the submodule courses for the operational officers (the language courses, the area studies courses, the professional studies courses), and the leadership courses for the senior officers (the senior leadership course, the ambassadorial seminar). The FSI's training pipeline is rank-progressive and role-specific, with the courses sequenced to match the officer's progression through the foreign service ranks. The FSI is the exemplar of the rank-progressive, role-specific training pipeline, and the geospatial training pipeline's design draws from the FSI's model.

The Canada School of Public Service (CSPS) is the training institution for the Canadian federal public service, established in 2004 (replacing the former Training and Development Canada and the Canadian Centre for Management Development). The CSPS offers the foundational courses (the orientation for the new public servants), the operational courses (the functional courses for the operational public servants), and the leadership courses (the Direction course for the executives, the Management course for the managers). The CSPS's training pipeline is government-wide, with the courses delivered to all federal public servants, and the CSPS is the exemplar of the government-wide training authority. The geospatial training pipeline's partnership with the administrative training authority draws from the CSPS's model, with the geospatial office as the client and the authority as the provider.

The French École nationale d'administration (ENA) is the training institution for the French senior civil service, established in 1945 and located in Strasbourg (since 1991). The ENA offers the rigorous curriculum for the senior civil service, with the competitive entry (the concours) and the 18-month curriculum (the public administration, the law, the economics, the international affairs, the leadership). The ENA's training pipeline is the elite training model, with the competitive entry and the rigorous curriculum, and the graduates are appointed to the senior civil service positions. The ENA is

the exemplar of the elite training model, and the geospatial training pipeline's rigor (the rigorous assessment, the competitive certification) draws from the ENA's model. The ENA was reformed in 2022 into the Institut national du service public (INSP), with the reform maintaining the elite training model while broadening the access.

The German Bundesakademie für öffentliche Verwaltung (BAköV) is the training institution for the German federal public service, established in 1969 and located in Brühl, North Rhine-Westphalia. The BAköV offers the foundational courses for the entry-level federal public servants, the operational courses for the operational public servants, and the leadership courses for the senior federal public servants. The BAköV's training pipeline is the federal training model, with the courses delivered to the federal and the state public servants, and the coordination with the state training institutions (the Landesakademien) through the joint committees. The BAköV is the exemplar of the federal training model, and the geospatial training pipeline's coordination with the sub-national geospatial offices draws from the BAköV's model.

The Singapore Civil Service College (CSC) is the training institution for the Singapore public service, established in 2001 (replacing the former Civil Service Institute). The CSC offers the foundational courses for the entry-level public servants (the Foundation Course), the operational courses for the operational public servants (the operational and the specialist courses), and the leadership courses for the senior public servants (the Strategic Leadership Programme, the Management Leadership Programme). The CSC's training pipeline is the comprehensive training model, with the courses delivered to all public servants across the service, and the CSC's close partnership with the public service's ministries and statutory boards. The CSC is the exemplar of the comprehensive training model, and the geospatial training pipeline's comprehensive coverage (the foundational, the submodule, the domain, the leadership) draws from the CSC's model.

The five models share the rank-progressive and role-specific design, and the five models are the models that the geospatial training pipeline draws from. The five models also share the certification's rigor and the continuing education's requirement, with the certifications maintained through the ongoing training and the officers' competence maintained over the career. The geospatial training pipeline's design is the synthesis of the five models, with the rank-progressive and role-specific structure (from the FSI), the government-wide partnership (from the CSPS), the rigorous assessment (from the ENA/INSP), the federal coordination (from the BAköV), and the comprehensive coverage (from the CSC). The synthesis is the geospatial training pipeline's distinctive design, and the design is the subject of this chapter.

31.9 Worked Example: The Training Pipeline of a New Resource Economics Specialist

The training pipeline is illustrated through the example of a new resource economics specialist who joins the office in January 2025 and completes the training pipeline through the foundational, the submodule, and the domain stages by December 2025. The example demonstrates the training pipeline's operation across the four stages and the certification's award and maintenance, and the example is the illustration of the training pipeline's institutional framework.

The specialist, Mr. Jonas Adekoya, is recruited in November 2024 and joins the office on 6 January 2025. Mr. Adekoya holds a Master of Science in Agricultural Economics from the University

of Northfields (2024), with the thesis on the crop yield forecasting using the satellite-based vegetation indices. He holds the WOFOST crop model certification (obtained in 2023) and the ASPRS Certified GIS/LIS Technologist certification (obtained in 2024). He has the reading proficiency in English (the international language) at the C1 level. His qualifications meet the entry requirements for the resource economics specialist role, and his recruitment is treated in the office's recruitment log.

The foundational stage begins on 13 January 2025, with Mr. Adekoya attending the foundational course at the administrative training authority's training centre. The course is delivered over 4 weeks (the 20-day course, with the 5 modules of 4 days each), with the morning sessions covering the theory and the afternoon sessions covering the practical exercises. The introduction to the geospatial profession module (week 1) covers the profession's history, the institutions, and the standards, with Mr. Adekoya's prior knowledge from his university studies supplemented by the office-specific content. The introduction to the pipeline architecture module (week 1) covers the pipeline's six capabilities, the shared-backbone and swappable-heads design, and the four domains, with Mr. Adekoya's prior knowledge of the food security domain supplemented by the other three domains' introduction. The introduction to the information system module (week 2) covers the six submodules, the integration with the core system, and the technology architecture. The introduction to the classification framework module (week 3) covers the four classification levels, the application, the legal framework, and the enforcement. The introduction to the office's procedures module (week 4) covers the daily, the event-driven, the commercial, and the reporting procedures. The foundational assessment is conducted on 7 February 2025, with the written test (the theory) and the practical test (the procedures). Mr. Adekoya passes the assessment with the score of 87 per cent, and the foundational certification is awarded on 10 February 2025.

The submodule stage begins on 17 February 2025, with Mr. Adekoya attending the modeling workbench course (the submodule most relevant to the resource economics specialist role). The course is delivered over 5 days, with the morning sessions covering the workbench environment, the domain plug-ins, the parameter management, the validation framework, and the provenance recording, and the afternoon sessions covering the practical exercises. The course's assessment is conducted on 21 February 2025, with the written test and the practical test (the modeling workbench operation, including the plug-in execution, the parameter configuration, the validation, the provenance recording). Mr. Adekoya passes the assessment with the score of 82 per cent, and the modeling workbench certification is awarded on 24 February 2025. Mr. Adekoya also attends the catalog and archive course (the submodule relevant to the specialist's product archiving) in March 2025, with the course delivered over 5 days and the assessment passed with the score of 85 per cent. The catalog and archive certification is awarded on 14 March 2025.

The domain stage begins on 17 March 2025, with Mr. Adekoya attending the crop model course (the domain course for the food security domain, his primary domain). The course is delivered over 10 days (the 5-day vegetation and phenology module and the 5-day yield forecasting and supply-demand module), with the morning sessions covering the theory and the afternoon sessions covering the practical exercises. The vegetation and phenology module (week 1) covers the NDVI, the EVI, the NDWI, the computation from the Sentinel-2 surface reflectance, the crop calendar, the accumulated growing degree days, and the stage transitions. The yield forecasting and supply-demand module (week 2) covers the statistical regression, the mechanistic models (the WOFOST and

the DSSAT), the calibration, the validation, the production estimate, the consumption estimate, the balance, and the release or buy signal. The course's assessment is conducted on 28 March 2025, with the written test and the practical test (the crop modelling, including the vegetation index computation, the phenological stage determination, the yield forecasting, the supply-demand analysis, the validation). Mr. Adekoya passes the assessment with the score of 91 per cent (the highest score in his cohort, reflecting his prior knowledge from his university studies), and the crop model certification is awarded on 31 March 2025.

The domain stage continues with the prospectivity course (the domain course for the minerals domain, his secondary domain) in September 2025. The course is delivered over 10 days (the 5-day prospectivity module and the 5-day resource estimation and validation module), with the morning sessions covering the theory and the afternoon sessions covering the practical exercises. The prospectivity module covers the knowledge-driven prospectivity (the boolean logic, the fuzzy logic, the weights of evidence), the data-driven prospectivity (the logistic regression, the random forest, the neural networks), and the geological data analysis. The resource estimation and validation module covers the kriging, the inverse distance weighting, the resource classification, and the validation protocols. The course's assessment is conducted on 12 September 2025, with the written test and the practical test. Mr. Adekoya passes the assessment with the score of 78 per cent, and the prospectivity certification is awarded on 15 September 2025.

The on-the-job training is conducted throughout the training pipeline, with Mr. Adekoya working alongside the senior resource economics specialist (the specialist treated in Chapter 18, section 18.6) on the office's standing production. In February-March 2025, Mr. Adekoya shadows the senior specialist on the weekly crop condition bulletin production, with the shadowing providing the practical experience of the bulletin's production. In April-May 2025, Mr. Adekoya takes the lead on the bulletin's production, with the senior specialist reviewing and validating the work. In June-August 2025, Mr. Adekoya works on the Northfields wheat yield forecasting (the forecasting that builds on the drought assessment treated in Chapter 25, section 25.9, with the 2025 season's recovery from the 2024 drought), with the work providing the practical experience of the yield forecasting and the supply-demand analysis. In September-October 2025, Mr. Adekoya works on the commercial minerals commission for the Westbourne Minerals Company (the prospectivity map for the Westbourne exploration area), with the work providing the practical experience of the commercial minerals intelligence production under the commercial firewall.

The continuing education begins in 2026, with Mr. Adekoya's training plan including the update course on the new Sentinel-2 Level-2A processing (the new processing baseline introduced by the ESA in late 2025, with the update course delivered in February 2026), the advanced course on the machine learning for the crop yield forecasting (the advanced course delivered in May 2026, with the course covering the random forest, the gradient boosting, and the neural networks for the yield forecasting), and the conference attendance at the ESA Living Planet Symposium 2026 (the conference in June 2026, with the attendance funded by the office's training budget and the presentation of the paper on the Northfields wheat yield forecasting). The continuing education is tracked through Mr. Adekoya's training record, with the record specifying the courses completed, the certifications maintained, and the continuing education hours accumulated.

The certification's maintenance begins in 2028, with Mr. Adekoya's crop model certification due for the recertification in March 2028 (the 3-year recertification cycle from the March 2025 award). The recertification requires the completion of the crop model refresher course (the 2-day course delivered in February 2028) and the accumulation of the 120 continuing education hours over the 3-year cycle (the hours accumulated through the update courses, the advanced courses, the conference attendance, the publication, and the professional service). Mr. Adekoya's recertification is verified by the office manager (who tracks the certifications) and ratified by the qualifications authority, and the crop model certification is recertified in March 2028.

This worked example illustrates the training pipeline's full operation: the foundational stage (the 4-week course in January-February 2025, the foundational certification on 10 February 2025), the submodule stage (the modeling workbench course in February 2025, the catalog and archive course in March 2025, the submodule certifications on 24 February and 14 March 2025), the domain stage (the crop model course in March 2025, the prospectivity course in September 2025, the domain certifications on 31 March and 15 September 2025), the on-the-job training (the shadowing in February-March, the lead production in April-May, the Northfields yield forecasting in June-August, the Westbourne commercial commission in September-October), the continuing education (the update course in February 2026, the advanced course in May 2026, the conference in June 2026), and the certification's maintenance (the recertification in March 2028). The training pipeline's institutional framework operated as designed: the foundational course established the common knowledge, the submodule courses built the technical competence, the domain courses built the analytical competence, the on-the-job training provided the practical experience, the continuing education maintained and updated the competence, and the certification's maintenance ensured the competence's currency. The framework is the condition of the office's ability to develop the officers' competence, and the framework is the reason the training pipeline exists. Mr. Adekoya's training pipeline, from the foundational course in January 2025 to the recertification in March 2028, is the illustration of the training pipeline's operation, and the pipeline is the example that this chapter provides for the administrative training authority's delivery and the qualifications authority's ratification. The training pipeline that is delivered and ratified becomes a permanent feature of the office's professional development, and the delivery and the ratification are the condition of the pipeline's sustainability, and the sustainability is the condition of the office's ability to maintain its competence over the long term.

Chapter 32 — Job Methods and Continuous Improvement

32.1 The Job Methods Tradition

Job methods and continuous improvement is the Training Within Industry Job Methods dimension applied to the geospatial office. The TWI Job Methods method was developed during the Second World War as part of the Training Within Industry programme, established by the United States War Manpower Commission in 1940 to support the wartime production expansion. The method was developed by the TWI Service under the leadership of Channing Rice Dooley, Walter Dietz, and Mike Kane, and it was based on the earlier work methods improvement approaches developed by Frank Gilbreth and Charles Bedaux in the early 20th century. The method was adopted by over 16,000 U.S. companies during the war, and after the war, the method was transmitted to Japan through the U.S. occupation's Civil Communications Section, where it became the foundation of the Toyota Production System and the broader Japanese continuous-improvement tradition.

The TWI Job Methods method is the method by which the worker improves the work method through the systematic questioning of every detail of the method. The method comprises four steps: break down the existing method, question every detail, develop the new method, and apply the new method. The first step (the breakdown) decomposes the existing method into its constituent details, with each detail recorded as a step, a key point, and a reason for the key point (the same structure used in the TWI Job Breakdown Sheets treated in Part III). The second step (the questioning) examines each detail and asks: What is done? Why is it done? Where is it done? When is it done? Who does it? How is it done? The questioning identifies the details that can be eliminated, simplified, combined, or rearranged. The third step (the development) develops the new method by applying the eliminations, the simplifications, the combinations, and the rearrangements identified in the questioning. The fourth step (the application) applies the new method, with the application documented, trained, and monitored.

The Job Methods tradition is the tradition that treats the work method as something that can and should be improved, and the tradition is the tradition that the geospatial office adopts for its own methods. An office that adopts the Job Methods tradition treats its methods as improvable, with the officers continuously questioning the methods and proposing the improvements. An office that does not adopt the tradition treats its methods as fixed, with the methods unchanged until they fail and the failure forces the change. The tradition is the institutional commitment to the continuous improvement of the office's methods, and the commitment is the condition of the office's resilience.

The Job Methods tradition is particularly important for a small geospatial office, because the small office's methods must be efficient and effective for the office to deliver its mandate with the limited staff. The small office cannot afford the cost of the suboptimal methods, because the small office has no slack to compensate: every seat's inefficiency is borne by the whole office, and the cumulative inefficiency undermines the office's ability to produce the intelligence its consumers require. The Job Methods tradition is the mechanism by which the small office identifies and eliminates the inefficiencies, and the tradition is the condition of the small office's ability to deliver its mandate with the limited staff.

32.2 The Method-Improvement Proposal

The method-improvement proposal is the process by which an officer proposes an improvement to an existing method, and it is the operational mechanism of the continuous-improvement discipline. The proposal comprises the description of the existing method, the identification of the improvement opportunity, the description of the proposed method, the justification for the improvement, and the assessment of the improvement's impact. The proposal is submitted to the technical lead, who reviews the proposal and either authorizes the improvement, requests further development, or declines the proposal with a documented reason. The proposal is the mechanism by which the officers' ideas for improvement are captured and evaluated, and the proposal is the mechanism by which the office's methods evolve.

The proposal's structure has five sections, each serving a specific purpose in the improvement's evaluation. The first section (the description of the existing method) documents the current method's steps, key points, and reasons, using the TWI Job Breakdown Sheet format treated in Part III. The description is the baseline against which the improvement is measured, and the description's accuracy is the condition of the improvement's evaluation. The second section (the identification of the improvement opportunity) identifies the specific detail or details of the existing method that the proposal improves, with the identification based on the Job Methods questioning (What? Why? Where? When? Who? How?). The identification is the analysis that justifies the improvement, and the identification's rigour is the condition of the improvement's credibility. The third section (the description of the proposed method) documents the new method's steps, key points, and reasons, using the same Job Breakdown Sheet format. The description is the specification of the improvement, and the description's clarity is the condition of the improvement's implementation. The fourth section (the justification) explains why the proposed method is better than the existing method, with the justification based on the elimination, the simplification, the combination, or the rearrangement of the details. The justification is the argument for the improvement, and the justification's persuasiveness is the condition of the improvement's authorization. The fifth section (the impact assessment) estimates the improvement's impact on the office's work, with the impact measured in the time saved, the quality improved, the risk reduced, or the cost avoided. The impact assessment is the quantification of the improvement's benefit, and the assessment's accuracy is the condition of the improvement's prioritization.

The proposal's review is conducted by the technical lead, with the review assessing the proposal's accuracy, rigour, clarity, persuasiveness, and accuracy. The technical lead's review may authorize the improvement (the improvement is implemented as proposed), request the further development (the improvement is returned to the officer for the revision), or decline the proposal (the improvement is not implemented, with the documented reason). The review's outcome is communicated to the officer within 10 business days of the proposal's submission, and the outcome is documented in the office's improvement record. The review's transparency is the condition of the officers' continued engagement with the proposal process, and the transparency is the mechanism that prevents the proposals from being silently ignored.

The proposal is not a suggestion box; the proposal is a structured process that requires the officer to think through the improvement, the justification, and the impact. The structured process is the mechanism that ensures the proposals are serious and evaluable, and the process is the mechanism

that prevents the proposals from being either too vague to evaluate or too numerous to manage. The office that maintains a structured proposal process produces improvements that are well-considered and well-implemented. The office that maintains an unstructured process produces improvements that are poorly considered and poorly implemented, with the poor implementation undermining the officers' confidence in the improvement process and the officers' engagement with the continuous improvement.

The proposal's implementation is the final step of the improvement process, with the implementation comprising the method's documentation (the new Job Breakdown Sheet), the training (the officers who operate the method are trained on the new method), the application (the new method is applied in the daily work), and the monitoring (the new method's performance is monitored to verify the improvement's impact). The implementation is the mechanism that translates the proposal into the actual improvement, and the implementation's completeness is the condition of the improvement's realization. The implementation is tracked through the improvement record, with the record specifying the implementation's timeline, the training's completion, and the monitoring's results.

32.3 The Annual Procedure Review

The annual procedure review is the process by which the office reviews its procedures on an annual basis, and it is the mechanism by which the office ensures its procedures remain current. The review comprises the identification of the procedures that need review, the review of each identified procedure, the revision of the procedures that need revision, and the documentation of the review. The review is conducted by the technical lead, with the officers who operate the procedures providing the input, and the office manager approving the revisions. The review is the mechanism by which the office's procedures are maintained, and the review is the condition of the procedures' currency.

The identification of the procedures that need review is the first step of the annual review, with the identification based on the triggers that warrant the review. The triggers comprise the method-improvement proposals that have been implemented (the procedures must be updated to reflect the improvements), the post-incident reviews' corrective actions (the procedures must be updated to reflect the lessons learned), the changes in the office's work (the procedures must be updated to reflect the new products, the new consumers, or the new domains), the changes in the information system (the procedures must be updated to reflect the new submodules, the new parameters, or the new configurations), and the changes in the external environment (the procedures must be updated to reflect the new regulations, the new standards, or the new inter-agency frameworks). The identification is conducted by the technical lead, with the office manager and the officers contributing the triggers' identification.

The review of each identified procedure is the second step, with the review examining the procedure's accuracy, completeness, clarity, and currency. The accuracy review checks whether the procedure's steps, key points, and reasons are correct, with the accuracy verified against the actual operation. The completeness review checks whether the procedure covers all the steps required for the operation, with the completeness verified against the operational experience. The clarity review checks whether the procedure is understandable to the officers who operate it, with the clarity verified through the officers' feedback. The currency review checks whether the procedure reflects the current

tools, the current parameters, and the current configurations, with the currency verified against the information system's current state. The review is conducted by the technical lead, with the officers who operate the procedure providing the operational input.

The revision of the procedures that need revision is the third step, with the revision updating the procedure to reflect the review's findings. The revision may be minor (the correction of the errors, the clarification of the ambiguities, the update of the references) or major (the restructuring of the procedure, the addition of the new steps, the elimination of the obsolete steps). The revision is documented in the procedure's version history, with the version history recording the revision's date, the revision's author, the revision's changes, and the revision's rationale. The revised procedure is reviewed by the technical lead and approved by the office manager, with the approval documented in the procedure's approval record.

The documentation of the review is the fourth step, with the documentation recording the review's scope, the review's findings, the revisions made, and the review's participants. The documentation is maintained in the office's improvement record, with the record serving as the evidence of the review's conduct and the basis for the next year's review. The documentation is also the input to the national audit authority's annual audit, with the audit examining the review's conduct and the revisions' quality.

The annual procedure review is not a formality; the review is a substantive examination of the procedures, and the review is the mechanism by which the procedures are updated to reflect the changes in the office's work, the information system, and the external environment. The review's substantive character is the condition of the procedures' currency and effectiveness, and the substantive character is the mechanism that prevents the procedures from becoming the outdated documents that the officers ignore.

32.4 The Post-Incident Review

The post-incident review is the process by which the office reviews its response to an incident after the incident is resolved, and it is the mechanism by which the office learns from the incident. The review comprises the timeline of the incident, the products produced, the issues encountered, the lessons learned, and the corrective actions. The review is conducted by the office manager, with the hazard specialist and the technical lead providing the input, and the review is documented. The review is the mechanism by which the office's incident response improves over time, and the review is the condition of the office's ability to respond better to the next incident.

The timeline of the incident is the chronological reconstruction of the incident's events, with the timeline recording the event's onset, the office's activation, the products produced, the products delivered, the authority's response, the event's peak, the event's recession, and the office's stand-down. The timeline is reconstructed from the office's event log, the production records, the communication records, and the authority's response records. The timeline's accuracy is the condition of the review's effectiveness, and the accuracy is verified through the cross-referencing of the multiple sources.

The products produced are the office's products during the incident, with the products comprising the event-specific intelligence (the inundation maps, the exposure reports, the evacuation guidance, treated in Chapter 24, section 24.4) and the event-driven updates (the 6-to-12-hour updates

through the event's evolution). The products' review examines the products' accuracy (the predictions compared with the observed data), the products' timeliness (the products' delivery compared with the authority's needs), and the products' usefulness (the products' use by the authority's response). The products' review is the mechanism that identifies the products' strengths and the weaknesses, and the review is the basis for the products' improvement.

The issues encountered are the problems that the office faced during the incident, with the issues comprising the data gaps (the data that the office needed but did not have), the model limitations (the model's inability to represent the event's specific characteristics), the communication failures (the communication channels that failed or were delayed), and the coordination gaps (the coordination issues with the authority, the water authority, or the other agencies). The issues' review examines each issue's cause, the issue's impact, and the issue's resolution. The issues' review is the mechanism that identifies the office's operational weaknesses, and the review is the basis for the corrective actions.

The lessons learned are the review's findings, with the lessons comprising the lessons about the office's methods (what the office did well and what it did poorly), the lessons about the office's tools (what the tools did well and what they did poorly), and the lessons about the office's coordination (what the coordination did well and what it did poorly). The lessons are the review's substantive output, and the lessons are the basis for the corrective actions. The lessons' candour is the condition of the corrective actions' effectiveness, and the candour is the mechanism that prevents the review from becoming the defensive exercise that protects the office's reputation at the expense of the learning.

The corrective actions are the actions the office takes to address the lessons learned, with the actions comprising the method improvements (the procedure revisions, the new procedures), the tool improvements (the parameter adjustments, the new configurations, the new tools), and the coordination improvements (the framework revisions, the new communication channels, the new coordination protocols). The corrective actions are documented in the office's improvement record, with the record specifying the action, the action's owner, the action's timeline, and the action's completion. The corrective actions are tracked through the monthly coordination meetings until the completion, and the completion is verified by the office manager.

The post-incident review is the most important continuous-improvement mechanism for the flood risk domain, because the flood events are infrequent and the office's event response must improve from event to event. An office that conducts the post-incident review honestly, acknowledging its failures as well as its successes, improves its event response. An office that does not, or that conducts the review defensively, does not improve, and the next event will reproduce the same failures. The post-incident review is the mechanism that makes the event response a learning process, and the learning allows the office to improve its event response over time, despite the infrequency of the events.

32.5 The Discipline of Not Freezing

The discipline of not freezing methods is the discipline of maintaining the methods as improvable, and it is the discipline that the Job Methods tradition requires. An office that freezes its methods treats the methods as fixed and unchangeable, with the methods unchanged until they fail and the failure forces the change. An office that does not freeze its methods treats the methods as

improvable, with the methods continuously questioned and improved. The discipline is the institutional commitment to the methods' continuous improvement, and the commitment is the condition of the office's resilience.

The freezing of the methods is the failure that the discipline is designed to prevent, and the freezing occurs for several reasons. The freezing may occur because the office's leadership is satisfied with the current methods and does not see the need for the improvement. The freezing may occur because the office's officers are too busy with the daily work to question the methods. The freezing may occur because the office's culture discourages the questioning of the established methods, with the questioning perceived as the criticism of the methods' authors. The freezing may occur because the office's information system locks in the methods, with the methods embedded in the system's configuration and the reconfiguration too costly to undertake. The discipline of not freezing is the mechanism that prevents each of these freezing causes, and the discipline is maintained through the leadership's commitment, the officers' time allocation, the culture's encouragement, and the system's flexibility.

The discipline of not freezing is particularly important for a small office, because the small office cannot afford the cost of the suboptimal methods. The large office can absorb the cost of the suboptimal methods, because the large office has the slack to compensate. The small office cannot, because the small office has no slack, and the cost of the suboptimal method is borne by every seat simultaneously. The small office that freezes its methods produces work that is increasingly inefficient, with the inefficiency accumulating over time and eventually undermining the office's ability to deliver its mandate. The discipline of not freezing is the mechanism that prevents the inefficiency's accumulation, and the discipline is the condition of the small office's resilience.

The discipline of not freezing is maintained through the institutional mechanisms that the Job Methods tradition provides. The method-improvement proposal (treated in section 32.2) is the mechanism that captures the officers' improvement ideas, with the proposal's structured process ensuring the ideas are serious and evaluable. The annual procedure review (treated in section 32.3) is the mechanism that ensures the procedures are reviewed and updated regularly, with the review's substantive character ensuring the procedures remain current. The post-incident review (treated in section 32.4) is the mechanism that ensures the incident response is reviewed and improved, with the review's candour ensuring the learning. The improvement record (treated in section 32.7) is the mechanism that documents the improvements, with the record's completeness ensuring the improvements are tracked and the audit-ready. The four mechanisms together constitute the institutional infrastructure for the discipline of not freezing, and the infrastructure is the condition of the discipline's maintenance.

The discipline of not freezing is not the absence of standards; the discipline is the maintenance of the standards as improvable. The office that has no standards has no methods to improve, and the office that has the frozen standards has the methods that cannot be improved. The office that has the improvable standards has the methods that can be improved, and the improvement is the mechanism that maintains the standards' currency and effectiveness. The discipline of not freezing is the discipline of the improvable standards, and the discipline is the synthesis of the standards' stability and the methods' improvement.

32.6 The Lean and Kaizen Traditions

The Lean and Kaizen traditions are the traditions that descend from the TWI Job Methods method, and they are the traditions that the geospatial office's continuous-improvement discipline draws from. The Lean tradition, developed by Toyota and elaborated by the Lean community, is the tradition of the systematic elimination of waste through the continuous improvement of the work method. The Kaizen tradition, developed in Japan and elaborated by the Kaizen community, is the tradition of the continuous, incremental improvement of the work method by the workers who perform the method. The two traditions share the Job Methods foundation, and the two traditions are the traditions that the geospatial office's continuous-improvement discipline is aligned with.

The Lean tradition originated in the Toyota Production System, developed by Taiichi Ohno and Shigeo Shingo at the Toyota Motor Company from the 1950s onward. Ohno, who joined Toyota in 1932 and became the production manager in the 1950s, developed the Toyota Production System as the response to the Japanese automotive industry's challenges: the small market, the limited capital, and the need for the high-variety, low-volume production. Ohno's system comprised the just-in-time production (the production of the right items at the right time in the right quantity), the jidoka (the autonomation, or the automation with the human touch, where the machines stop automatically when the abnormality is detected), and the kanban (the visual signalling system that controls the production flow). Ohno's system was elaborated by Shingo, who contributed the poka-yoke (the mistake-proofing, or the design of the process to prevent the errors) and the single-minute exchange of die (the SMED, or the reduction of the setup time). The Lean tradition was transmitted to the West in the 1980s, with the publication of "The Machine That Changed the World" by James Womack, Daniel Jones, and Daniel Roos in 1990, and the Lean tradition was adopted by the manufacturing, the service, and the government sectors.

The Kaizen tradition originated in the Japanese post-war industrial reconstruction, with the Kaizen concept elaborated by Masaaki Imai in his 1986 book "Kaizen: The Key to Japan's Competitive Success." Imai defined the Kaizen as the continuous, incremental improvement of the work method by the workers who perform the method, with the improvement driven by the workers' ideas rather than the management's directives. The Kaizen tradition comprises the Kaizen events (the short, focused improvement events where the workers identify and implement the improvements), the suggestion processes (the structured processes for the workers to submit the improvement ideas), and the worker-led improvement (the improvement driven by the workers' ownership of the methods). The Kaizen tradition is the tradition that the geospatial office's method-improvement proposal (treated in section 32.2) is most closely aligned with, with the proposal as the geospatial office's variant of the Kaizen suggestion process.

The Lean and Kaizen traditions provide the tools and the techniques that the geospatial office's continuous-improvement discipline uses. The Lean tradition provides the value-stream mapping (the mapping of the work's flow from the input to the output, with the value-adding and the non-value-adding steps identified), the waste identification (the identification of the seven wastes: the overproduction, the waiting, the transport, the over-processing, the inventory, the motion, and the defects), and the root-cause analysis (the analysis of the problem's root cause, using the 5 Whys or the fishbone diagram). The Kaizen tradition provides the continuous improvement events (the short, focused events where the workers identify and implement the improvements), the suggestion

processes (the structured processes for the workers to submit the improvement ideas), and the worker-led improvement (the improvement driven by the workers' ownership of the methods).

The geospatial office adapts these tools and techniques to the government geospatial context, and the adaptation is the mechanism by which the traditions are made applicable to the office's work. The value-stream mapping is adapted to the geospatial office's pipeline (the mapping of the data's flow from the acquisition through the archive, with the value-adding and the non-value-adding steps identified). The waste identification is adapted to the geospatial office's work (the identification of the geospatial wastes: the over-processing, the redundant processing, the unnecessary format conversions, the manual operations that could be automated, the errors that require the rework, the waiting for the data, the motion of the data between the systems). The root-cause analysis is adapted to the geospatial office's problems (the analysis of the model's error, the processing's failure, or the product's defect, using the 5 Whys to identify the root cause). The continuous improvement events are adapted to the geospatial office's schedule (the half-day or the full-day events where the officers identify and implement the improvements, with the events scheduled quarterly). The suggestion processes are adapted to the geospatial office's method-improvement proposal (treated in section 32.2). The worker-led improvement is adapted to the geospatial office's culture (the officers' ownership of their methods, with the officers responsible for the methods' improvement).

The adaptation is not the uncritical adoption of the traditions; the adaptation is the thoughtful application of the traditions' principles to the office's context. The geospatial office's context differs from the manufacturing context where the Lean tradition originated, and the differences require the adaptation. The geospatial office's work is the analytical work, not the production work, and the analytical work's improvement is the improvement of the analytical methods, not the improvement of the production flow. The geospatial office's products are the intelligence products, not the physical products, and the intelligence products' quality is the analytical quality, not the manufacturing quality. The geospatial office's consumers are the government agencies, not the market customers, and the government agencies' needs are the institutional needs, not the market preferences. The adaptation is the mechanism by which the traditions' principles are applied to the geospatial office's context, and the adaptation is the mechanism by which the traditions contribute to the office's continuous improvement.

32.7 The Improvement Record

The improvement record is the record of the improvements the office has made to its methods, and it is the record that allows the office to track its continuous-improvement progress. The record comprises the method-improvement proposals, the annual procedure reviews, the post-incident reviews, and the implemented improvements. The record is maintained by the technical lead, with the office manager reviewing, and the record is the institutional memory of the office's continuous-improvement efforts. The record is the mechanism by which the office's improvements are documented, and the record is the condition of the office's ability to learn from its improvement efforts.

The improvement record's structure has four sections, each corresponding to a type of the improvement. The method-improvement proposals section records each proposal's submission date, the proposing officer, the proposal's content, the technical lead's review, the review's outcome, and (if

authorized) the implementation's timeline and completion. The annual procedure reviews section records each review's scope, the review's findings, the revisions made, and the review's participants. The post-incident reviews section records each incident's timeline, the products produced, the issues encountered, the lessons learned, and the corrective actions. The implemented improvements section records each improvement's implementation date, the implementation's owner, the implementation's monitoring results, and the improvement's impact.

The improvement record's maintenance is the technical lead's responsibility, with the record updated as the improvements are proposed, reviewed, implemented, and monitored. The record is stored in the office's information system, with the access restricted to the office's officers and the audit authority. The record's retention is permanent, with the improvements retained for the office's institutional memory. The record's format is the structured document (the spreadsheet or the database), with the format allowing the record's analysis (the improvements' frequency, the improvements' impact, the improvements' trends).

The improvement record is also the record that the national audit authority examines during the annual audit, because the record is the evidence of the office's continuous-improvement discipline. The audit authority's examination checks whether the office has the continuous-improvement discipline (the record's existence), whether the discipline is substantive (the record's content), and whether the discipline is effective (the record's impact). The audit authority's examination is the external accountability for the office's continuous-improvement discipline, and the examination is the mechanism that ensures the discipline's maintenance.

The improvement record is also the input to the office's annual reporting, with the annual report documenting the office's continuous-improvement progress (the number of the proposals submitted, the number of the proposals implemented, the impact of the implemented improvements, the lessons learned from the post-incident reviews). The annual report's continuous-improvement section is the office's public account of its improvement efforts, and the section is the mechanism by which the office demonstrates its commitment to the continuous improvement to the desk chief, the bureau chief, and the audit authority.

The improvement record is the institutional mechanism for the office's continuous-improvement accountability, and the mechanism is the condition of the office's ability to demonstrate its discipline to the audit authority. The office that maintains the improvement record produces the evidence of its discipline. The office that does not maintain the record produces no evidence, and the absence of evidence is the failure that the record is designed to prevent. The improvement record is the institutional expression of the office's commitment to the continuous improvement of its methods, and the record is the condition of the office's resilience.

32.8 The Comparative Practice

The comparative practice of the continuous-improvement discipline is examined through the institutions that have adopted the TWI Job Methods tradition and the Lean and Kaizen traditions. The comparative practice spans the manufacturing sector (where the traditions originated), the government sector (where the traditions have been adapted), and the geospatial sector (where the traditions are applied to the geospatial work). The comparative practice is examined in four

institutions, with the four institutions representing the variants of the continuous-improvement discipline's application.

The Toyota Motor Company is the institution where the Lean tradition originated, with the Toyota Production System developed by Taiichi Ohno and Shigeo Shingo from the 1950s onward. Toyota's continuous-improvement discipline comprises the Kaizen Teian (the suggestion system, where the workers submit the improvement ideas and the ideas are evaluated and implemented), the Kaizen Events (the short, focused events where the workers identify and implement the improvements), and the Hansei (the reflection, where the workers reflect on their work and identify the improvements). Toyota's discipline is the exemplar of the continuous-improvement discipline's manufacturing application, and the geospatial office's method-improvement proposal draws from the Toyota's Kaizen Teian. The Toyota Production System's documentation (Ohno's "Toyota Production System: Beyond Large-Scale Production" in 1988, Shingo's "A Study of the Toyota Production System" in 1989) is the reference for the Lean tradition's principles and the tools.

The Virginia Mason Medical Center is the institution where the Lean tradition has been adapted to the healthcare sector, with the Virginia Mason Production System developed by the Virginia Mason Medical Center in Seattle, Washington, from 2002 onward, in partnership with the Lean Enterprise Institute. The Virginia Mason Production System comprises the Rapid Process Improvement Workshops (the week-long workshops where the teams identify and implement the improvements), the Production System Daily Management (the daily management of the work using the Lean principles), and the Patient Safety Alert System (the mistake-proofing system where the staff report the errors and the near-misses). The Virginia Mason Production System is the exemplar of the continuous-improvement discipline's service-sector application, and the geospatial office's event-driven improvement draws from the Virginia Mason's Rapid Process Improvement Workshops.

The United States federal government is the institution where the Lean tradition has been adapted to the government sector, with the Lean Government initiative promoted by the U.S. Environmental Protection Agency (EPA) and the National Institute of Standards and Technology (NIST) from the 2000s onward. The Lean Government initiative comprises the Lean Kaizen Events (the events where the government teams identify and implement the improvements in the government processes), the Value-Stream Mapping (the mapping of the government processes' flow), and the Lean Six Sigma (the combination of the Lean's waste elimination with the Six Sigma's defect reduction). The Lean Government initiative is the exemplar of the continuous-improvement discipline's government-sector application, and the geospatial office's annual procedure review draws from the Lean Government's Value-Stream Mapping.

The European Space Agency (ESA) is the institution where the continuous-improvement discipline has been applied to the geospatial work, with the ESA's Earth Observation Programme Board's continuous-improvement framework established in the 2010s. The ESA's framework comprises the Mission Performance Review (the annual review of the satellite missions' performance, with the review identifying the improvements), the Product Quality Improvement (the ongoing improvement of the satellite products' quality, with the improvements driven by the user feedback), and the Calibration and Validation Campaigns (the regular campaigns to calibrate and validate the satellite data, with the campaigns' findings driving the improvements). The ESA's

framework is the exemplar of the continuous-improvement discipline's geospatial application, and the geospatial office's post-incident review draws from the ESA's Mission Performance Review.

The four institutions share the TWI Job Methods foundation, with the four institutions' continuous-improvement disciplines all derived from the Job Methods' four steps (the breakdown, the questioning, the development, the application). The four institutions also share the worker-led improvement character, with the four institutions' improvements all driven by the workers who perform the methods, rather than the management's directives. The four institutions differ in the context (the manufacturing, the healthcare, the government, the geospatial), but the continuous-improvement discipline's principles are the same across the contexts. The geospatial office's continuous-improvement discipline is the synthesis of the four institutions' disciplines, with the method-improvement proposal (from the Toyota's Kaizen Teian), the event-driven improvement (from the Virginia Mason's Rapid Process Improvement Workshops), the annual procedure review (from the Lean Government's Value-Stream Mapping), and the post-incident review (from the ESA's Mission Performance Review). The synthesis is the geospatial office's distinctive continuous-improvement discipline, and the discipline is the subject of this chapter.

32.9 Worked Example: The Post-Incident Review of the Riverbourne Basin Flood Event

The post-incident review is illustrated through the Riverbourne Basin flood event of March 2024, treated in Chapters 24, 27, 28, and 29. The event's response involved the office's flood risk intelligence production, the coordination with the emergency management authority (treated in Chapter 24), the coordination with the water authority (treated in Chapter 27), and the operation of the geospatial information system (treated in Chapter 29). The post-incident review is the mechanism by which the office learns from the event and improves its methods, tools, and coordination for the next event.

The review is conducted on 5 April 2024, twelve days after the office's stand-down on 28 March 2024. The review is conducted by the office manager, with the hazard specialist (Dr. Maren Holloway, treated in Chapter 30, section 30.9), the technical lead, the data pipeline operator, and the production operator providing the input. The review is conducted in the office's conference room, with the review's duration of 4 hours (the morning session covering the timeline and the products, the afternoon session covering the issues and the lessons). The review's documentation is prepared by the office manager, with the documentation comprising the timeline, the products, the issues, the lessons, and the corrective actions.

The timeline of the incident is reconstructed from the office's event log, the production records, and the communication records. The timeline records: the meteorological service's severe weather warning at 09:15 on 23 March 2024; the office's notification of the emergency management authority at 09:30; the office's activation at 09:45; the accelerated waterway data sharing with the Riverbourne Water Authority at 10:30; the HEC-RAS model run at 11:00; the event-specific flood inundation map at 12:30; the product delivery to the emergency management authority at 13:30 and to the water authority at 13:00; the evacuation orders at 14:30; the office manager's liaison deployment at 16:00; the 6-hour update cycle through 23-26 March; the event's peak at 04:00 on 25 March (the river stage of 4.2 metres above the flood stage); the revised inundation map at 06:00 on 25 March; the event's

recession on 26-28 March; and the office's stand-down on 28 March. The timeline's reconstruction is verified through the cross-referencing of the event log, the production records, and the communication records, with the verification confirming the timeline's accuracy.

The products produced during the incident are reviewed, with the products comprising the event-specific flood inundation map (the no-breach and the breach scenarios, produced at 12:30 on 23 March), the event-specific exposure report (produced at 13:00), the event-specific evacuation guidance (produced at 13:00), the 6-hour updates through 23-26 March (12 updates in total), and the revised inundation map at 06:00 on 25 March. The products' review examines the products' accuracy (the no-breach scenario's inundation extent was 12 square kilometres, the observed extent was 18 square kilometres, an underestimation of 33 per cent due to the underestimation of the storm surge; the breach scenario's inundation extent was 18 square kilometres, the observed extent was 18 square kilometres, an exact match but the breach did not occur, and the match is coincidental rather than predictive), the products' timeliness (the first products were delivered within 4 hours of the activation, meeting the 2-to-4-hour target), and the products' usefulness (the products were used by the emergency management authority for the evacuation orders, the resource deployment, and the public warnings, confirming the products' operational value).

The issues encountered during the incident are reviewed, with the issues comprising: (1) the underestimation of the no-breach scenario's inundation extent, caused by the underestimation of the storm surge's contribution (the model's storm surge boundary condition used the deterministic forecast, which underestimated the observed storm surge by 20 per cent); (2) the Levee Section 23's settlement, which was known to the water authority but not communicated to the office in the standing waterway data sharing (the settlement was communicated only in the event-driven data sharing, when the office asked for the levee conditions); (3) the 6-hour update cycle's strain on the office's surge capacity, with the hazard specialist and the production operator working 16-hour shifts through the event's duration; (4) the communication gap between the office's liaison officer (the office manager, deployed to the emergency management authority's emergency operations centre) and the operational responders, with the operational responders occasionally producing their own intelligence in parallel with the office's products.

The lessons learned from the event are the review's substantive findings, with the lessons comprising: (1) the storm surge boundary condition must be improved, with the use of the ensemble forecast rather than the deterministic forecast, to capture the storm surge's uncertainty and to produce the inundation scenarios that reflect the uncertainty's range; (2) the waterway data sharing must include the levee settlement data in the standing data sharing, not only in the event-driven data sharing, so that the office's flood models can incorporate the settlement in the standing model runs; (3) the office's surge capacity must be increased, with the cross-training of the resource economics specialist on the hazard specialist's tasks, to provide the hazard specialist with the relief during the extended events; (4) the communication with the operational responders must be improved, with the office's liaison officer providing the direct interpretation of the products to the operational responders, and the operational responders' parallel intelligence production discouraged through the emergency management authority's common operating picture's enforcement.

The corrective actions are the actions the office takes to address the lessons learned, with the actions comprising: (1) the HEC-RAS model's storm surge boundary condition is updated to use the

ensemble forecast, with the update implemented by the hazard specialist by 1 June 2024 and verified by the technical lead by 15 June 2024; (2) the data-sharing agreement with the Riverbourne Water Authority is revised to include the levee settlement data in the standing data sharing, with the revision negotiated by the office manager by 1 May 2024 and ratified by the water authority by 1 June 2024; (3) the resource economics specialist (Mr. Jonas Adekoya, treated in Chapter 31, section 31.9) is cross-trained on the hazard specialist's tasks, with the cross-training conducted by the hazard specialist in May-June 2024 and the cross-training's completion verified by the technical lead by 1 July 2024; (4) the liaison officer's role is formalized in the inter-agency framework with the emergency management authority, with the role's responsibilities and the communication protocols documented by the office manager by 1 June 2024 and agreed by the emergency management authority by 1 July 2024.

The corrective actions are documented in the office's improvement record, with the record specifying the action, the action's owner, the action's timeline, and the action's completion. The corrective actions are tracked through the monthly coordination meetings, with the office manager reviewing the actions' progress and the technical lead reporting the actions' implementation. The corrective actions' completion is verified by the office manager, with the verification documented in the improvement record. The corrective actions are completed by 1 July 2024, and the office's improved methods, tools, and coordination are operational for the 2025 flood season.

The post-incident review's documentation is maintained in the office's improvement record, with the documentation comprising the timeline, the products, the issues, the lessons, and the corrective actions. The documentation is the input to the national audit authority's annual audit, with the audit examining the review's conduct and the corrective actions' implementation. The documentation is also the input to the next post-incident review, with the next review comparing the next event's response with this event's response and assessing the corrective actions' effectiveness.

This worked example illustrates the continuous-improvement discipline's full operation: the post-incident review's conduct (the 4-hour review on 5 April 2024, with the office manager and the officers), the timeline's reconstruction (the event's chronology from the meteorological warning at 09:15 on 23 March to the stand-down on 28 March), the products' review (the accuracy, the timeliness, the usefulness), the issues' identification (the storm surge underestimation, the levee settlement data gap, the surge capacity strain, the communication gap), the lessons' derivation (the four lessons corresponding to the four issues), the corrective actions' development (the four actions with the owners and the timelines), the corrective actions' implementation (the completion by 1 July 2024), and the documentation's maintenance (the improvement record, the audit input, the next review's input). The continuous-improvement discipline's institutional framework operated as designed: the post-incident review was conducted honestly (the failures were acknowledged, including the 33-per-cent underestimation and the communication gap), the lessons were derived rigorously (the four lessons corresponding to the four issues, with the lessons' candour reflecting the review's candour), the corrective actions were implemented effectively (the four actions completed by 1 July 2024, with the completion verified), and the documentation was maintained (the improvement record, the audit input, the next review's input). The framework is the condition of the office's ability to learn from its events and to improve its event response over time, despite the infrequency of the major events. The Riverbourne Basin flood event demonstrated the continuous-

improvement discipline's value: the post-incident review identified the office's failures (the storm surge underestimation, the levee settlement data gap, the surge capacity strain, the communication gap), the corrective actions addressed the failures, and the office's 2025 flood season's response is improved by the corrective actions. The continuous-improvement discipline is the mechanism by which the office learns from its events, and the learning is the reason the discipline exists, and the discipline is the condition of the office's resilience, the office's effectiveness, and the office's contribution to the public safety during the flood events.

Chapter 33 — Comparative Geospatial Intelligence Traditions

Part VII — Comparative Practice and Craft: The comparative traditions, the officer's judgment, drafting and communication, and the toolkit.

33.1 The Purpose of Comparative Study

Comparative geospatial intelligence traditions are examined through six national institutions, each representing a distinct model of how a government organizes its geospatial intelligence function. The purpose of the comparative study is not to identify a single best model but to understand the range of models that comparative practice has produced, the conditions under which each model is appropriate, and the lessons that each model offers. The geospatial office that understands the comparative traditions can draw from the full range of practice; an office that does not is confined to its own tradition, with the confinement limiting the office's ability to address the challenges that its own tradition does not adequately address.

The six institutions are the United States Geological Survey (USGS), the European Space Agency's Copernicus programme, the French Institut national de l'information géographique et forestière (IGN), Geoscience Australia (GA), Indonesia's Badan Informasi Geospasial (BIG), and the Canada Centre for Mapping and Earth Observation (CCMEO). Each institution is examined for the model it represents, the history that shaped it, the functions it performs, the products it delivers, and the lessons it offers. The examination is conducted with attention to what each model contributes to the discipline and where each model's limits lie, because no model is without its limits, and the honest assessment of the limits is the condition of the comparative study's usefulness.

The comparative study's method is the structured comparison, with each institution examined against the same analytical framework. The framework comprises six dimensions: the institutional history (the institution's establishment, evolution, and current form), the functional scope (the functions the institution performs), the funding model (the public appropriation, the commercial revenue, the cost recovery), the data policy (the free and open, the restricted, the commercial), the organizational placement (the ministry, the agency, the statutory authority), and the lessons offered (the contributions to the discipline and the limits). The six dimensions are the analytical lenses through which each institution is examined, and the dimensions are the mechanism by which the institutions are compared.

The comparative study is the institutional foundation for the geospatial office's doctrine, treated throughout this Handbook. The commercial firewall doctrine (treated in Chapter 5) draws from the USGS and the GA. The free-and-open data policy (treated in Chapter 8) draws from the Copernicus programme. The cadastral authority interface (treated in Chapter 23) draws from the IGN. The compound hazard assessment (treated in Chapter 19) draws from the BIG. The national spatial data infrastructure (treated in Chapter 8 and Chapter 29) draws from the CCMEO. The doctrine's foundations in the comparative practice are the reason the doctrine is robust, and the foundations are the reason the doctrine is applicable across the jurisdictions that adopt it.

33.2 The United States Geological Survey

The United States Geological Survey (USGS) represents the minerals-funded public geoscience model, in which the commercial minerals function funds the public geoscience function through the revenue generated by mineral resource data sales and custom products. The Survey was established on 3 March 1879 by the U.S. Congress through the Sundry Civil Appropriations Act, with the initial mandate to classify the public lands and examine the geological structure, mineral resources, and products of the national domain. The Survey's first Director, Clarence King (1879-1881), established the Survey's scientific rigour, and King's successor John Wesley Powell (1881-1894) expanded the Survey's scope to include the topographic mapping and the hydrography. The Survey's mission has expanded over the 145 years since its establishment to include the hazards, resources, and environment programs that constitute the Survey's public geoscience portfolio.

The Survey's current organization comprises the Core Science Systems Mission Area (the National Geospatial Program, the National Land Imaging Program, the Science Synthesis, Analysis, and Research Program), the Ecosystems Mission Area, the Energy and Minerals Mission Area (the Mineral Resources Program, the Energy Resources Program), the Natural Hazards Mission Area (the Earthquake Hazards Program, the Volcano Hazards Program, the Landslide Hazards Program, the Global Seismographic Network), and the Water Mission Area (the National Water Quality Assessment Program, the National Water Information System, the Water Science Centers). The Survey operates from its headquarters in Reston, Virginia, and its regional centers: the Denver Federal Center in Lakewood, Colorado (the Earth Resources Observation and Science Center, EROS, in Sioux Falls, South Dakota; the National Geospatial Technical Operations Center in Denver; and the Water Resources Mission Area's western operations).

The Survey's Earth Resources Observation and Science (EROS) Center, established in 1971 in Sioux Falls, South Dakota, is the Survey's Earth observation data center. EROS operates the Landsat satellite ground system (treated in Chapter 8, section 8.2), archives and distributes the Landsat imagery (the Landsat archive comprising over 50 years of the Earth's land surface imagery, the longest continuous satellite imagery archive in the world), and produces the value-added products (the Landsat Analysis Ready Data, the Land Change Monitoring, Assessment, and Projection, the National Land Cover Database). EROS's archive also includes the Sentinel-2 and the Sentinel-1 imagery from the Copernicus programme, the MODIS imagery from the NASA Terra and Aqua satellites, and the high-resolution commercial imagery acquired through the National Geospatial-Intelligence Agency's EnhancedView programme.

The Survey's minerals-funded model is the model where the commercial minerals function funds the public geoscience function through the revenue generated by mineral resource data sales and custom products. The Survey's Mineral Resources Program, through the National Minerals Information Center (NMIC), compiles the mineral commodity summaries (the annual statistics on the U.S. and the global mineral production, consumption, and trade, published in the Minerals Yearbook and the Mineral Commodity Summaries), and the NMIC sells the value-added mineral intelligence to the commercial customers (the mining companies, the consulting firms, the financial institutions). The revenue from the mineral intelligence sales, combined with the appropriation from the Congress, funds the Survey's public programs. The model is the exemplar of the revenue-neutral geospatial

office, treated in Chapter 5, and the model is the model that the commercial firewall doctrine draws from.

The Survey's contribution to the discipline is the demonstration that the revenue-neutral model is viable at scale, and that the commercial minerals arm can sustain the public-service arm without compromising the arm's independence. The Survey has maintained the commercial firewall for over 145 years, with the firewall institutionalized in the Survey's procedures, the audit regime, and the professional culture. The Survey's limits lie in the model's dependence on the mineral sector's commercial demand, which fluctuates with the commodity cycle (the mineral intelligence sales revenue declines during the commodity downturns, and the decline pressures the public programs' funding). The dependence is the reason the smoothing mechanism (the revolving fund, treated in Chapter 5, section 5.3) is a necessary component of the revenue-neutral model. The Survey's lesson is that the revenue-neutral model works, but it works only with the firewall, the smoothing mechanism, and the audit regime that the model requires.

33.3 The Copernicus Programme

The European Space Agency's Copernicus programme represents the free-and-open data model, in which the Earth observation data is acquired by the Sentinel satellite constellation and distributed freely to all users, including the commercial sector, the research community, and the public. The programme was established in 2014 as the European Union's Earth observation programme, building on the earlier Global Monitoring for Environment and Security (GMES) initiative that began in 1998. The programme is the world's largest free-and-open Earth observation data source, with the Sentinel constellation comprising seven satellite missions (the Sentinel-1 radar, the Sentinel-2 optical, the Sentinel-3 ocean and land, the Sentinel-4 atmospheric, the Sentinel-5P atmospheric, the Sentinel-5 atmospheric, and the Sentinel-6 altimetry) and the contributing missions (the national and the commercial satellites that contribute the data to the programme).

The programme's governance is shared between the European Commission (the programme's overall direction and the funding), the European Space Agency (the satellite development and the operations), the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT, the Sentinel-3, 4, 5, and 6 operations), and the European Centre for Medium-Range Weather Forecasts (ECMWF, the atmosphere service). The programme's six services are the Land Monitoring Service (operated by the European Environment Agency and the Joint Research Centre), the Marine Environment Monitoring Service (operated by Mercator Ocean International), the Atmosphere Monitoring Service (operated by the ECMWF), the Emergency Management Service (operated by the Joint Research Centre, treated in Chapter 24, section 24.8), the Security Service (operated by the European Commission and the member states), and the Climate Change Service (operated by the ECMWF).

The programme's data policy is the free and open data policy, with the Sentinel data distributed freely to all users through the Copernicus Open Access Hub, the Copernicus Data and Information Access Services (DIAS, the cloud-based platforms that provide the data access and the processing), and the national mirror sites. The free and open data policy was established by the European Commission's Delegated Regulation 1159/2013, and the policy has transformed the European Earth observation market, with the Copernicus data generating an estimated 2 billion euros of economic

value annually and supporting over 17,000 jobs in the European Earth observation sector by 2020. The free and open data policy has also transformed the global Earth observation market, with the Copernicus data used by the government agencies, the commercial companies, and the research communities worldwide.

The Copernicus programme's contribution to the discipline is the demonstration that the free-and-open data model transforms the geospatial intelligence market, by making the base data available to all users at no cost, and by allowing the intelligence-producing offices to focus on the analysis rather than the data acquisition. The free-and-open data has lowered the barrier to entry for the geospatial intelligence production, with the small offices and the developing-country offices able to access the high-quality Earth observation data without the procurement costs. The free-and-open data has also accelerated the innovation, with the commercial companies and the research communities developing the new applications and the new methods on the Copernicus data.

The programme's limits lie in the model's dependence on the public funding of the satellite constellation, which is a substantial fiscal commitment (the programme's 2021-2027 budget is 5.4 billion euros, comprising the satellite development, the operations, and the services). The dependence is the reason the free-and-open model is not universally adopted, with many jurisdictions unable or unwilling to make the fiscal commitment. The programme's lesson is that the free-and-open data model is transformative where it is feasible, and that the geospatial office that has access to the free-and-open data should use it as the foundation for its intelligence production, and should reserve its commercial imagery budget for the applications that the free-and-open data cannot serve (the high-resolution monitoring, the sub-daily revisit, the tasking).

33.4 The French Institut national de l'information géographique et forestière

The French Institut national de l'information géographique et forestière (IGN) represents the cadastre-integrated mapping model, in which the cadastre, the land registry, and the national mapping function are combined in a single institution. The IGN was established in its current form in 2012 through the merger of the Institut géographique national (the national mapping agency, established in 1940) and the Inventaire forestier national (the national forest inventory, established in 1958). The IGN operates under the Ministry of Ecology (the Ministère de la Transition écologique et de la Cohésion des territoires), with the head office in Saint-Mandé, near Paris, and the regional offices across the French territory.

The IGN's functions comprise the cadastre (maintained in partnership with the Direction générale des finances publiques, DGFiP, the tax authority that holds the legal authority for the cadastre), the topographic mapping (the 1:25,000 topographic map series, the SCAN 25, and the 1:100,000 and the 1:250,000 map series), the administrative boundaries (the official boundaries of the communes, the departments, and the regions), the address database (the Base Adresse Nationale, the national address database), the orthophoto (the BD ORTHO, the national orthophoto coverage at 50-centimetre resolution), the elevation data (the RGE Alti, the national elevation database at 1-metre resolution from the LiDAR), and the forest inventory (the national forest inventory, with the field plots and the remote sensing data).

The IGN's data is distributed through the IGN's geospatial portal (the Géoportail, established in 2006 as the public access to the IGN's data) and the standardized web services (the OGC WMS, WFS,

and WCS services). The data policy is the free and open data policy for the public-sector users and the commercial license for the commercial users, with the data policy established by the French government's open data policy (the Loi pour une République numérique, the Digital Republic Act of 2016). The IGN's data is the base layer for the French government's analytical agencies, including the Institut national de la statistique et des études économiques (INSEE, the national statistics agency, treated in the comparative practice of Chapter 28), the Bureau de recherches géologiques et minières (BRGM, the geological survey), and the agricultural research institutes.

The IGN's contribution to the discipline is the demonstration that the cadastre-integrated model produces consistent, high-quality base data, because the cadastre and the mapping are maintained by the same institution and the data consistency is internal. The IGN's base data is among the highest-quality national base data in the world, with the cadastral accuracy, the topographic accuracy, and the elevation accuracy all maintained to the high standards. The IGN's limits lie in the model's concentration of institutional power, which may create its own conflicts of interest (the IGN's dominance in the French geospatial market may limit the competition and the innovation), and the concentration is the reason the integrated model is not universally adopted. The IGN's lesson is that the cadastre-integrated model is effective where the jurisdiction has chosen it, and that the geospatial office that operates under the integrated model benefits from the data consistency, but must maintain the analytical independence that the integrated model's institutional power may threaten.

The IGN's history reflects the French tradition of the centralized state and the scientific cartography, with the IGN's predecessor institutions dating to the 18th century (the Académie des Sciences' Cassini map of France, the first national topographic survey, begun in 1748 and completed in 1815). The IGN's current form, with the 2012 merger, reflects the modern French government's effort to consolidate the geospatial functions for the efficiency and the data consistency. The IGN's organization, with the scientific staff and the technical staff organized in the engineers and the technicians corps (treated in Chapter 30, section 30.8), reflects the French civil service's professional model, and the organization is the model that the geospatial role family's qualifications framework draws from.

33.5 Geoscience Australia

Geoscience Australia (GA) represents the cost-recovery model applied to a national geological survey, in which the commercial minerals intelligence sales fund the public geoscience programs. The agency was established in its current form on 1 July 2001 through the merger of the Australian Geological Survey Organisation (AGSO, the former Bureau of Mineral Resources, established in 1946) and the Australian Surveying and Land Information Group (AUSLIG, the former Division of National Mapping, established in 1947). The merger consolidated the geoscience, the geological mapping, the hazard assessment, the geodesy, and the remote sensing functions in a single agency. The GA operates under the Minister for Resources, with the head office in Canberra.

The GA's functions comprise the mineral prospectivity (the mineral prospectivity models, the geological maps, the mineral resource assessments, treated in Chapter 22), the petroleum prospectivity (the offshore petroleum prospectivity, the sedimentary basin analysis), the hazard assessment (the earthquake hazard, the flood hazard, the tsunami hazard, the volcanic hazard, the landslide hazard), the geodesy (the Australian Regional GPS Network, the geodetic datum

maintenance, the National Positioning Infrastructure), the Earth observation (the Digital Earth Australia programme, the satellite imagery archive, the Landsat and the Sentinel data processing), and the spatial data infrastructure (the Australian Spatial Data Infrastructure, the Place Names Register, the administrative boundaries).

The GA's Digital Earth Australia (DEA) programme, established in 2018, is the GA's Earth observation programme, providing the free and open access to the Australian Landsat and Sentinel data through the cloud-based platform. The DEA programme is the GA's variant of the Copernicus free-and-open data model, with the DEA providing the data access and the processing platform to the government, the research, and the commercial users. The DEA programme has transformed the Australian Earth observation market, with the DEA data used for the agricultural monitoring, the water monitoring, the coastal monitoring, and the bushfire monitoring.

The GA's cost-recovery model is the model where the commercial minerals intelligence sales fund the public geoscience programs. The GA's Minerals Division sells the mineral prospectivity models, the geological maps, and the resource estimates to the mining and exploration companies, with the sales revenue contributing to the GA's budget. The model is combined with the public appropriation from the Australian government, with the appropriation funding the public programs (the hazard assessment, the geodesy, the Earth observation) and the cost-recovery funding the commercial programs (the minerals intelligence). The model is the exemplar of the minerals-commercial model in a jurisdiction with a substantial mineral sector (Australia is the world's second-largest mineral exporter, after China), and the model is the model that the geospatial office's commercial minerals arm draws from.

The GA's contribution to the discipline is the demonstration that the cost-recovery model is viable in a jurisdiction with a strong mineral sector, and that the commercial minerals arm can operate with the pricing flexibility and the customer orientation that the cost-recovery requires. The GA's Minerals Division has developed the customer-oriented products and the competitive pricing that the Australian mineral sector's commercial demand requires, and the Division's revenue has sustained the GA's minerals program through the commodity cycles. The GA's limits lie in the model's vulnerability to the commodity cycle, which affects the commercial demand (the mineral intelligence sales revenue declines during the commodity downturns, with the 2015-2016 commodity downturn reducing the Division's revenue by approximately 30 per cent), and the vulnerability is the reason the revolving fund (treated in Chapter 5, section 5.3) is a necessary component. The GA's lesson is that the cost-recovery model works where the mineral sector is strong, and that the geospatial office that adopts the model must maintain the firewall and the revolving fund that the model requires.

The GA's organization, with the scientific staff organized in the Australian Public Service classification (treated in Chapter 30, section 30.8), and the GA's partnership with the state geological surveys (the state geological surveys that regulate the mineral sector, treated in Chapter 26, section 26.8), reflect the Australian federal model's distribution of the geospatial functions across the national and the state governments. The GA's role as the national geoscience agency, with the state geological surveys as the sectoral regulators, is the model that the geospatial office's sectoral regulator interface draws from.

33.6 The Indonesian Badan Informasi Geospasial

The Indonesian Badan Informasi Geospasial (BIG) represents the flood-prone archipelagic state model, in which the geospatial intelligence function is organized to serve the compound flood hazards of an archipelagic state. The BIG was established in its current form on 4 April 2011 by the Indonesian Law No. 4 of 2011 on Geospatial Information, replacing the former Badan Koordinasi Survei dan Pemetaan Nasional (Bakosurtanal, the National Coordinating Agency for Surveys and Mapping, established in 1969). The BIG operates under the President of Indonesia, with the head office in Cibinong, West Java, and the regional offices across the Indonesian archipelago. The BIG's establishment as a non-ministerial government agency reporting directly to the President reflects the Indonesian government's recognition of the geospatial intelligence's strategic importance.

The BIG's functions comprise the national spatial data infrastructure (the Indonesian National Spatial Data Infrastructure, the SDI, with the standardized base data and the metadata), the topographic mapping (the Indonesia Topographic Map, the RBI, at 1:25,000, 1:50,000, and 1:250,000 scales), the bathymetric mapping (the Indonesian Bathymetric Map, with the hydrographic surveys of the Indonesian waters), the thematic mapping (the land cover, the land use, the administrative boundaries), the hazard mapping (the flood hazard, the volcano hazard, the earthquake hazard, the tsunami hazard, the landslide hazard), the geodesy (the Indonesian Geospatial Reference System, the geodetic datum maintenance), and the geospatial information services (the data distribution through the BIG's portal and the web services).

The BIG's flood hazard assessment is the function that exemplifies the flood-prone archipelagic state model. Indonesia is one of the world's most flood-prone countries, with the compound flood hazards comprising the riverine flooding (from the seasonal monsoon rainfall, affecting the major river basins in Sumatra, Java, Kalimantan, and Sulawesi), the coastal flooding (from the high tides and the sea-level rise, affecting the low-lying coastal areas where the majority of the Indonesian population lives), and the storm surge (from the tropical cyclones that affect the eastern Indonesian islands). The BIG's flood hazard assessment integrates the riverine, the coastal, and the storm surge hazards, producing the compound flood hazard maps that the Indonesian National Disaster Management Authority (BNPB, the Badan Nasional Penanggulangan Bencana) uses for the disaster preparedness and the response.

The BIG's contribution to the discipline is the demonstration that the geospatial intelligence function can serve the most demanding flood hazards, and that the compound hazard assessment (the integration of the multiple flood hazards into a single assessment) is a capability that the geospatial office can develop. The BIG's compound flood hazard assessment is among the most sophisticated in the world, with the integration of the hydrological modelling (the HEC-HMS and the HEC-RAS models), the coastal modelling (the MIKE 21 model), and the storm surge modelling (the ADCIRC model), and the assessment's methodology is the model that the geospatial office's flood risk intelligence (treated in Chapter 19) draws from.

The BIG's limits lie in the resource constraints that affect the geospatial intelligence function in many jurisdictions. Indonesia, despite being the world's fourth-most populous country and the largest archipelagic state, has the geospatial intelligence budget that is modest compared to the U.S. or the European budgets, and the modest budget limits the BIG's ability to acquire the high-resolution commercial imagery, to maintain the dense monitoring networks, and to develop the advanced

analytical capabilities. The resource constraints are the reason the revenue-neutral model (treated in Chapter 5) may be necessary even in the jurisdictions where the mineral sector is not strong, and the constraints are the reason the free-and-open data (the Copernicus data, the Landsat data) is the critical foundation for the BIG's intelligence production. The BIG's lesson is that the flood-prone state model is a model that any jurisdiction with significant flood hazards can draw from, and that the compound hazard assessment is a capability that the geospatial office should develop regardless of its jurisdiction's geography.

33.7 The Canada Centre for Mapping and Earth Observation

The Canada Centre for Mapping and Earth Observation (CCMEO) represents the national spatial data infrastructure model, in which the geospatial intelligence function is organized around the national spatial data infrastructure that coordinates the production and the sharing of spatial data across the government. The CCMEO was established in its current form through the reorganization of the national mapping function within Natural Resources Canada (NRCan), with the CCMEO operating under the Department of Natural Resources. The CCMEO's head office is in Ottawa, with the Earth observation data center in Ottawa and the regional offices across Canada.

The CCMEO's functions comprise the national topographic base (the Canadian National Topographic System, the NTS, at 1:50,000 and 1:250,000 scales), the national spatial data infrastructure (the Canadian Geospatial Data Infrastructure, the CGDI, with the standardized data and the web services), the Earth observation data archive (the Canada Centre for Remote Sensing, the CCRS, established in 1971, with the satellite imagery archive including the Landsat, the Sentinel, the RADARSAT, and the MODIS data), the geodesy (the Canadian Active Control System, the Canadian Spatial Reference System), and the federal geospatial coordination (the Federal Geospatial Platform, the FGP, with the federal geospatial data and the services).

The CCMEO's Federal Geospatial Platform (FGP), established in 2011, is the CCMEO's coordination mechanism for the federal geospatial activities. The FGP comprises the federal geospatial data catalogue (the discovery and the access to the federal geospatial data), the federal geospatial web services (the standardized web services for the federal geospatial data), and the federal geospatial governance (the inter-departmental coordination on the geospatial standards, the data sharing, and the joint projects). The FGP's participants include the CCMEO (the topographic base, the Earth observation, the geodesy), the Statistics Canada (the census geography, treated in Chapter 28), the Environment and Climate Change Canada (the weather, the climate, the hydrology), the Department of Fisheries and Oceans (the hydrography, the maritime boundaries), the Agriculture and Agri-Food Canada (the agricultural land use), and the other federal departments with the geospatial data.

The CCMEO's data policy is the free and open data policy for the public-sector users, established by the Canadian government's Open Data initiative (launched in 2011) and the Open Government License (established in 2013). The CCMEO's data is distributed through the Open Maps portal (the Canadian government's geospatial open data portal) and the standardized web services. The CCMEO's RADARSAT data (the Canadian SAR satellite constellation, the RADARSAT-1 launched in 1995, the RADARSAT-2 launched in 2007, and the RADARSAT Constellation Mission launched

in 2019) is the commercial data, with the data distribution through the MDA Corporation (the satellite operator) and the limited free access for the government users.

The CCMEO's contribution to the discipline is the demonstration that the national spatial data infrastructure model is a viable alternative to the sole-producer model, and that the coordination of the distributed spatial data production can produce consistent, high-quality spatial data without concentrating the production in a single institution. The CCMEO's FGP has enabled the federal departments to share the geospatial data, to avoid the duplication, and to coordinate the standards, with the FGP's coordination reducing the federal geospatial expenditures and improving the data's consistency. The CCMEO's limits lie in the coordination costs that the distributed model incurs (the inter-departmental coordination requires the significant effort, with the FGP's governance, the standards development, and the data sharing agreements), and the costs are the reason the sole-producer model remains attractive for the jurisdictions that can afford it. The CCMEO's lesson is that the national spatial data infrastructure model is effective where the jurisdiction has chosen it, and that the geospatial office that operates under the model must invest in the coordination mechanisms that the model requires.

The CCMEO's history reflects the Canadian federal model's distribution of the geospatial functions across the federal and the provincial governments, with the CCMEO as the federal geospatial agency and the provincial geospatial agencies (the provincial mapping agencies, the provincial land registries) as the provincial counterparts. The CCMEO's partnership with the provincial agencies, through the Canadian Geomatics Accord (the inter-governmental accord on the geospatial coordination) and the Canadian Council on Geomatics (the inter-governmental council), is the model that the geospatial office's inter-agency interfaces draw from.

33.8 The Lessons of Comparative Practice

The lessons of comparative practice are the lessons that the six institutions collectively offer, and the lessons are the lessons that the geospatial office draws from the comparative study. The six lessons are the substantive findings of the comparative analysis, and the lessons are the institutional foundations for the geospatial office's doctrine.

The first lesson, from the USGS, is that the revenue-neutral model is viable at scale, but only with the commercial firewall, the smoothing mechanism, and the audit regime that the model requires. The USGS has maintained the revenue-neutral model for over 145 years, with the firewall institutionalized in the procedures, the audit, and the culture. The geospatial office that adopts the revenue-neutral model must adopt the firewall, the smoothing mechanism (the revolving fund), and the audit regime, and the office that adopts the model without the requirements will find the model unsustainable. This lesson is the foundation for the commercial firewall doctrine treated in Chapter 5.

The second lesson, from the Copernicus programme, is that the free-and-open data model transforms the geospatial intelligence market, and the office that has access to the free-and-open data should use it as the foundation for its intelligence production. The Copernicus programme's free and open data has lowered the barrier to entry, accelerated the innovation, and allowed the intelligence-producing offices to focus on the analysis rather than the data acquisition. The geospatial office that has access to the Copernicus data (or the Landsat data, or the other free-and-open data) should use the data as the foundation, and should reserve its commercial imagery budget for the applications that the

free-and-open data cannot serve. This lesson is the foundation for the data acquisition strategy treated in Chapter 8.

The third lesson, from the IGN, is that the cadastre-integrated model produces consistent, high-quality base data, but the office that operates under the model must maintain its analytical independence. The IGN's cadastre-integrated model has produced the base data that is among the highest-quality in the world, with the consistency maintained internally. The geospatial office that operates under the integrated model benefits from the data consistency, but the office must maintain the analytical independence that the integrated model's institutional power may threaten. This lesson is the foundation for the cadastral authority interface treated in Chapter 23.

The fourth lesson, from the GA, is that the cost-recovery model works where the mineral sector is strong, and the office that adopts the model must maintain the firewall and the revolving fund. The GA's cost-recovery model has sustained the GA's minerals program through the commodity cycles, with the revolving fund smoothing the revenue fluctuations. The geospatial office that adopts the cost-recovery model in a jurisdiction with a strong mineral sector must maintain the firewall and the revolving fund, and the office that adopts the model in a jurisdiction without a strong mineral sector will find the model unsustainable. This lesson is the foundation for the commercial minerals intelligence treated in Chapter 22 and the commercial firewall treated in Chapter 5.

The fifth lesson, from the BIG, is that the flood-prone state model is a model that any jurisdiction with significant flood hazards can draw from, and the compound hazard assessment is a capability that the office should develop. The BIG's compound flood hazard assessment, integrating the riverine, the coastal, and the storm surge hazards, is among the most sophisticated in the world, and the assessment's methodology is applicable to any jurisdiction with the compound flood hazards. The geospatial office that serves a flood-prone jurisdiction should develop the compound hazard assessment capability, regardless of the jurisdiction's geography. This lesson is the foundation for the flood risk intelligence treated in Chapter 19.

The sixth lesson, from the CCMEO, is that the national spatial data infrastructure model is effective where the jurisdiction has chosen it, and the office that operates under the model must invest in the coordination mechanisms. The CCMEO's FGP has enabled the federal departments to share the geospatial data and to coordinate the standards, with the coordination reducing the expenditures and improving the consistency. The geospatial office that operates under the distributed model must invest in the coordination mechanisms (the governance, the standards, the data sharing agreements), and the office that operates under the sole-producer model benefits from the simplicity but incurs the full cost of the production. This lesson is the foundation for the spatial data infrastructure treated in Chapter 8 and the geospatial information system specification treated in Chapter 29.

The six lessons are the lessons that the comparative practice offers, and the lessons are the lessons that the geospatial office's doctrine, treated throughout this Handbook, is built on. The doctrine's foundations in the comparative practice are the reason the doctrine is robust, and the foundations are the reason the doctrine is applicable across the jurisdictions that adopt it. The geospatial office that understands the comparative traditions can draw from the full range of practice, and the office that draws from the full range of practice produces the doctrine that is informed by the world's best models and the doctrine that is adapted to the office's own context.

33.9 Worked Example: A Comparative Analysis for a New Geospatial Office

The comparative analysis is illustrated through the example of a new geospatial office established in 2025 by the government of a fictional jurisdiction, the Republic of Maritimea, with the office's establishment requiring the comparative analysis to inform the office's design. Maritimea is a mid-sized jurisdiction (approximately 15 million population, approximately 250,000 square kilometres of territory) with the compound flood hazards (the riverine flooding, the coastal flooding), a moderate mineral sector (the copper, the gold, the lithium), and the developing agricultural sector (the wheat, the rice, the maize). The Maritimean government's decision to establish the geospatial office is driven by the 2023 Riverbourne Basin flood event (treated in Chapters 24, 27, 28, 29, 32), which demonstrated the need for the geospatial intelligence to support the emergency management, the water management, and the food security.

The comparative analysis is conducted by the Maritimean government's establishment committee, with the committee comprising the Maritimean government's representatives (the Ministry of Planning, the Ministry of Finance, the Ministry of Science and Technology), the external advisors (the consultants from the international geospatial advisory firms), and the international observers (the representatives from the USGS, the Copernicus programme, the IGN, the GA, the BIG, and the CCMEO). The analysis is conducted over 6 months (the January-June 2025 period), with the analysis comprising the desk study (the review of the six institutions' documentation), the field visits (the visits to the six institutions' headquarters), and the synthesis (the comparative analysis and the recommendations).

The desk study is conducted in January-February 2025, with the committee reviewing the six institutions' documentation (the annual reports, the strategic plans, the operational plans, the data policies, the organizational charts, the budget documents). The desk study's output is the structured comparison of the six institutions against the six dimensions of the analytical framework (the institutional history, the functional scope, the funding model, the data policy, the organizational placement, and the lessons offered). The structured comparison identifies the models that are most relevant to the Maritimean context: the revenue-neutral model (the USGS and the GA), the free-and-open data model (the Copernicus programme), the cadastre-integrated model (the IGN), the flood-prone state model (the BIG), and the national spatial data infrastructure model (the CCMEO).

The field visits are conducted in March-May 2025, with the committee visiting the six institutions' headquarters (the USGS in Reston, Virginia; the Copernicus programme in Brussels, Belgium; the IGN in Saint-Mandé, France; the GA in Canberra, Australia; the BIG in Cibinong, Indonesia; the CCMEO in Ottawa, Canada). The field visits comprise the briefings by the institutions' senior staff, the tours of the institutions' facilities (the data centers, the processing facilities, the map libraries), and the discussions with the institutions' operational staff (the specialists, the operators, the managers). The field visits' output is the detailed understanding of each institution's operation, with the understanding going beyond the desk study's documentation to the operational reality that the documentation does not fully capture.

The synthesis is conducted in June 2025, with the committee producing the comparative analysis and the recommendations. The comparative analysis identifies the five models that are most relevant to the Maritimean context, and the analysis assesses each model's applicability to the Maritimean context. The revenue-neutral model (the USGS and the GA) is assessed as applicable, given the

Maritimean mineral sector's moderate strength, with the model's requirements (the firewall, the smoothing mechanism, the audit regime) identified as the prerequisites. The free-and-open data model (the Copernicus programme) is assessed as applicable, given the Maritimean government's ability to access the Copernicus data and the Landsat data, with the model's requirement (the commercial imagery budget for the high-resolution applications) identified. The cadastre-integrated model (the IGN) is assessed as not directly applicable, given the Maritimean cadastre's separate institutional location, but the model's lesson (the data consistency) is identified as the goal for the Maritimean cadastral authority interface. The flood-prone state model (the BIG) is assessed as highly applicable, given the Maritimean compound flood hazards, with the model's compound hazard assessment identified as the priority capability. The national spatial data infrastructure model (the CCMEO) is assessed as applicable, given the Maritimean government's distributed geospatial functions, with the model's coordination mechanisms identified as the requirement.

The recommendations are the committee's proposals for the Maritimean geospatial office's design, based on the comparative analysis. The first recommendation is that the Maritimean geospatial office should adopt the revenue-neutral model, with the commercial firewall, the smoothing mechanism, and the audit regime, drawing from the USGS and the GA. The second recommendation is that the office should adopt the free-and-open data policy, using the Copernicus and the Landsat data as the foundation, and reserving the commercial imagery budget for the high-resolution monitoring (treated in Chapter 21). The third recommendation is that the office should develop the compound flood hazard assessment capability, drawing from the BIG, with the assessment integrating the riverine, the coastal, and the storm surge hazards. The fourth recommendation is that the office should participate in the national spatial data infrastructure, drawing from the CCMEO, with the office as the coordinator of the Maritimean geospatial data sharing. The fifth recommendation is that the office should establish the cadastral authority interface, drawing from the IGN, with the interface ensuring the data consistency between the office's analytical products and the cadastre's legal record.

The recommendations are presented to the Maritimean government in July 2025, with the government's review and the approval. The government approves the recommendations, with the approval establishing the Maritimean Office of Geospatial Services (the MOGS) under the Ministry of Science and Technology, with the revenue-neutral model, the free-and-open data policy, the compound flood hazard assessment capability, the national spatial data infrastructure coordination, and the cadastral authority interface. The MOGS's establishment is the culmination of the comparative analysis, and the establishment is the institutional act that translates the comparative practice into the Maritimean context.

The MOGS's operation, from its establishment in 2025 onward, is the test of the comparative analysis's recommendations. The MOGS's first year (2026) is the year of the office's setup, with the recruitment of the officers (treated in Chapter 30), the training pipeline's establishment (treated in Chapter 31), the information system's procurement (treated in Chapter 29), and the inter-agency frameworks' negotiation (treated in Part V). The MOGS's second year (2027) is the year of the office's operational capability, with the first products' production (the flood hazard maps, the crop condition bulletins, the infrastructure monitoring, the mineral prospectivity), the first consumers' engagement (the emergency management authority, the strategic reserve authority, the sectoral regulator, the

water authority, the national statistics agency), and the first event response (the 2027 flood season). The MOGS's third year (2028) is the year of the office's institutional consolidation, with the first post-incident review (treated in Chapter 32), the first annual procedure review, and the first audit by the national audit authority.

This worked example illustrates the comparative study's full operation: the comparative analysis's purpose (the establishment committee's design of the Maritimean geospatial office), the desk study (the structured comparison of the six institutions against the six dimensions), the field visits (the briefings, the tours, the discussions), the synthesis (the comparative analysis and the recommendations), the recommendations' adoption (the Maritimean government's approval of the MOGS's establishment), and the MOGS's operation (the first three years' setup, operational capability, and institutional consolidation). The comparative study's institutional framework operated as designed: the six institutions were examined against the same analytical framework, the models' applicability to the Maritimean context was assessed, the recommendations were based on the comparative analysis, and the recommendations were adopted and implemented. The comparative study's value is the connection between the world's best models and the Maritimean office's design, and the connection is the reason the comparative study exists. The MOGS's establishment, based on the comparative analysis, is the illustration of the comparative study's operation, and the establishment is the example that this chapter provides for the geospatial office's design. The comparative study that is conducted seriously, with the desk study, the field visits, and the synthesis, produces the recommendations that are informed by the world's best models and adapted to the office's own context, and the recommendations are the condition of the office's effective establishment, and the office's effective establishment is the condition of the office's contribution to the jurisdiction's geospatial intelligence.

Chapter 34 — The Geospatial Officer's Judgment

34.1 The Nature of Professional Judgment

The geospatial officer's judgment is the judgment exercised in the situations where the rules do not determine the outcome, and the officer must weigh the competing considerations to reach a decision. The judgment is not the application of a rule but the weighing of competing considerations in the absence of a determinative rule, and it is the craft dimension of the geospatial profession. The rules, the procedures, and the job breakdown sheets treated throughout this Handbook cover the routine work, but the routine work is punctuated by the situations where the rules are silent, the procedures are ambiguous, or the job breakdown sheets do not apply, and the officer's judgment is the faculty that resolves these situations.

The professional judgment is distinguished from the rule application by three characteristics. The first characteristic is the ambiguity: the situation does not have a single correct answer that the rules can determine, but multiple plausible answers that the judgment must weigh. The second characteristic is the consequence: the decision has the consequences that are significant for the office's work, the consumers' decisions, or the public's safety, and the consequence is the reason the judgment matters. The third characteristic is the accountability: the officer is accountable for the decision, with the accountability documented in the office's records and the decision subject to the review by the senior officers, the audit authority, or the public inquiry. The three characteristics together distinguish the judgment from the rule application, and the characteristics are the conditions that make the judgment a professional faculty rather than a mechanical operation.

The professional judgment is developed through experience, through the study of comparative cases, and through the mentorship of the senior officers. The experience is the accumulation of the situations in which the judgment is exercised, with the officer encountering the situations, exercising the judgment, and learning from the encounter. The study is the examination of the comparative cases, through the after-action reviews, the professional literature, and the lessons-learned documents, with the officer learning from the experience of others. The mentorship is the guidance of the senior officers, who share their judgment with the junior officers through the working relationships, with the judgment transmitted across the generations. The three sources together constitute the judgment's development, and the development is the process by which the officer acquires the craft.

The judgment is the craft dimension of the geospatial profession, and the craft is the dimension that the training, the qualifications, and the experience are designed to develop. The training (treated in Chapter 31) provides the rules and the procedures, but the training does not provide the judgment, because the judgment cannot be taught in the classroom. The qualifications (treated in Chapter 30) verify the officer's knowledge and the competence, but the qualifications do not verify the judgment, because the judgment is not tested by the certification examination. The experience is the only source of the judgment, and the experience is the reason the judgment is the faculty of the experienced officer rather than the novice. The judgment's development over the career is the reason the geospatial profession is a career rather than a job, and the career is the condition of the judgment's acquisition.

34.2 When Models Disagree with Ground Truth

The first situation in which the geospatial officer's judgment is exercised is when the models disagree with the ground truth. The disagreement may arise from a model error, a data error, a validation error, or a genuine change in the conditions that the model does not capture. The officer must diagnose the cause of the disagreement, and the diagnosis requires the judgment to distinguish between the possible causes and to determine the appropriate response. An officer who diagnoses the disagreement as a model error when it is a genuine change will produce intelligence that misses the change; an officer who diagnoses the disagreement as a genuine change when it is a model error will produce intelligence that is wrong. The judgment is the faculty that distinguishes the two diagnoses, and the judgment is developed through the experience of the model-ground truth disagreements.

The model error is the disagreement caused by the model's inability to represent the physical process. The hydrological model may overestimate the runoff because the curve number is too high for the actual soil conditions; the hydraulic model may overestimate the inundation because the Manning's n value is too low for the actual channel roughness; the crop model may overestimate the yield because the crop coefficient is too high for the actual crop variety. The model error is diagnosed by the validation (the comparison of the model's predictions with the observed data) and the sensitivity analysis (the examination of the model's response to the parameter changes). The model error's response is the model's recalibration, with the parameters adjusted to match the observed data.

The data error is the disagreement caused by the error in the input data. The rainfall data may be underestimated because the rain gauge is sheltered by the nearby buildings; the river stage data may be incorrect because the gauge is reading against the wrong datum; the soil moisture data may be biased because the satellite's retrieval algorithm is not calibrated for the local conditions. The data error is diagnosed by the data quality checks (the comparison of the data with the independent sources) and the data validation (the examination of the data's consistency with the historical patterns). The data error's response is the data's correction, with the data replaced or adjusted based on the independent sources.

The validation error is the disagreement caused by the error in the validation data. The observed river stage may be incorrect because the gauge was overtopped during the flood; the observed inundation extent may be incomplete because the satellite imagery was acquired after the peak; the observed yield may be inaccurate because the field survey was conducted before the harvest. The validation error is diagnosed by the validation data's quality checks and the validation data's consistency with the other sources. The validation error's response is the validation data's correction or the validation's deferral until the better data is available.

The genuine change is the disagreement caused by the actual change in the conditions that the model does not capture. The channel geometry may have changed because of the sediment deposition or the bank erosion; the land cover may have changed because of the urbanization or the deforestation; the levee may have settled because of the foundation consolidation. The genuine change is diagnosed by the field verification (the physical inspection of the changed conditions) and the temporal analysis (the examination of the change's progression over time). The genuine change's response is the model's update, with the model's input data and the parameters updated to reflect the changed conditions.

The comparative cases of the model-ground truth disagreement are drawn from the published after-action reviews of the national geospatial agencies. The United States Geological Survey's flood model validations have encountered disagreements caused by the changes in the channel geometry that the model did not capture (the 2017 Hurricane Harvey after-action review documented the channel geometry changes in the Houston area's bayous that the Federal Emergency Management Agency's flood hazard maps did not reflect), and the Survey's response was to update the channel geometry and re-run the model. The European Flood Awareness System's flood forecasts have encountered disagreements caused by the differences in the rainfall data from the different sources (the 2021 European floods after-action review documented the differences between the radar rainfall estimates and the rain gauge observations), and the System's response was to compare the data sources and use the most reliable. The cases confirm that the model-ground truth disagreement is a recurring situation, and the judgment to diagnose and respond is a recurring exercise of the geospatial officer's craft.

34.3 When Data Is Ambiguous

The second situation in which the geospatial officer's judgment is exercised is when the data is ambiguous. The ambiguity may arise from the cloud cover that obscures the optical imagery, the radar artifacts that mimic the real features, the temporal mismatch between the data sources, or the resolution mismatch between the data and the analysis. The officer must assess the ambiguity, determine whether the data is sufficient for the analysis, and decide whether to proceed, to acquire additional data, or to defer the analysis. An officer who proceeds with insufficient data produces intelligence that is unreliable; an officer who defers unnecessarily produces intelligence that is late.

The cloud cover ambiguity is the ambiguity caused by the clouds that obscure the optical imagery. The cloud cover is the persistent challenge for the optical Earth observation, particularly in the tropical regions where the cloud cover is near-constant, and in the high-latitude regions where the cloud cover is frequent during the winter. The cloud cover ambiguity is assessed through the cloud mask (the identification of the cloud-covered pixels) and the cloud cover percentage (the ratio of the cloud-covered pixels to the total pixels). The officer must judge whether the cloud-free pixels are sufficient for the analysis (the analysis can proceed with the partial coverage), or whether the cloud cover requires the additional data (the radar imagery that penetrates the clouds) or the deferral (the wait for the next cloud-free acquisition).

The radar artifact ambiguity is the ambiguity caused by the artifacts in the radar imagery that mimic the real features. The radar artifacts include the layover (the radar returns from the tall features that appear to lean toward the sensor), the foreshortening (the radar returns from the slopes that appear compressed), the shadow (the radar returns from the areas behind the tall features that are not illuminated), and the speckle (the radar returns from the coherent illumination that produce the granular pattern). The officer must judge whether the artifact is a real feature (the change detection's identification of the real change) or an artifact (the false positive that should be excluded), and the judgment requires the understanding of the radar imaging geometry and the radar processing.

The temporal mismatch ambiguity is the ambiguity caused by the difference in the acquisition times of the data sources. The flood inundation map may combine the satellite imagery acquired at 10:00 with the hydraulic model's simulation at 06:00, and the four-hour difference may produce the

mismatch if the flood is evolving rapidly. The crop condition assessment may combine the Sentinel-2 imagery acquired on 15 July with the Sentinel-1 imagery acquired on 18 July, and the three-day difference may produce the mismatch if the crop is developing rapidly. The officer must judge whether the temporal mismatch is acceptable (the conditions are stable enough that the mismatch does not affect the analysis) or unacceptable (the conditions are changing too rapidly for the mismatch), and the judgment requires the understanding of the phenomenon's temporal dynamics.

The resolution mismatch ambiguity is the ambiguity caused by the difference in the spatial resolution of the data sources. The flood inundation map may combine the 10-metre Sentinel-2 imagery with the 2-metre LiDAR DEM, and the resolution difference may produce the mismatch if the flood's edge is at the sub-10-metre scale. The change detection may combine the 0.3-metre Maxar imagery with the 10-metre Sentinel-2 imagery, and the resolution difference may produce the mismatch if the change is at the sub-10-metre scale. The officer must judge whether the resolution mismatch is acceptable (the analysis can be conducted at the coarser resolution) or unacceptable (the analysis requires the finer resolution), and the judgment requires the understanding of the analysis's spatial requirements.

The comparative cases of data ambiguity are drawn from the professional literature of the national geospatial agencies. The crop condition assessments in the tropical regions, where the persistent cloud cover makes the optical imagery ambiguous, have been addressed by the fusion of the optical and the radar imagery (the Sentinel-2 and the Sentinel-1 fusion, treated in Chapter 9, section 9.6), and the fusion is the response to the ambiguity that the judgment produces. The change detection in the urban areas, where the shadow and the building height make the high-resolution imagery ambiguous, has been addressed by the use of the stereo imagery and the elevation data, and the use is the response to the ambiguity that the judgment produces. The cases confirm that the data ambiguity is a recurring situation, and the judgment to assess and respond is a recurring exercise of the geospatial officer's craft.

34.4 When the Firewall Is Tested

The third situation in which the geospatial officer's judgment is exercised is when the commercial firewall is tested by a customer who is also a government official. The situation arises when a senior official in a consumer ministry requests a commercial minerals product, or requests a modification to a public-service product that would benefit a commercial customer, and the pressure is to treat the request as a government request rather than a commercial request. The officer must recognize the test, apply the firewall's authority mechanism, and refuse to treat the request as a government request, regardless of the official's seniority. An officer who fails to recognize the test erodes the firewall; an officer who recognizes the test but yields to the pressure erodes the firewall more.

The firewall test takes several forms. The first form is the direct request: the senior official directly requests the commercial minerals product without the commercial charge code, arguing that the product is for the government's use and should be provided under the public-service appropriation. The second form is the indirect request: the senior official requests a "briefing" or a "summary" of the commercial minerals product, arguing that the briefing is not the product itself and does not require the commercial charge code. The third form is the modification request: the senior

official requests a modification to a public-service product (for example, the addition of the mineral prospectivity layer to the public geological map) that would benefit a commercial customer, arguing that the modification is the public-service product's enhancement and does not require the commercial charge code. The fourth form is the access request: the senior official requests the access to the commercial minerals archive, arguing that the access is for the government's oversight and does not require the commercial charge code. The four forms are the variants of the firewall test, and the officer must recognize each form and apply the firewall's authority mechanism.

The firewall's authority mechanism is the mechanism by which the officer refuses the test, and the mechanism is treated in Chapter 5, section 5.4. The mechanism comprises the charge code verification (every commercial product requires the commercial charge code, and the absence of the code is the firewall's trigger), the office manager's authorization (the commercial release requires the office manager's authorization, and the office manager is the firewall's authority), and the documentation (the firewall test is documented in the office's firewall log, with the documentation serving as the audit trail). The officer who encounters the firewall test must apply the mechanism: verify the charge code, refer the request to the office manager, and document the test. The mechanism is the institutional defence against the firewall test, and the mechanism is the condition of the firewall's integrity.

The firewall test's recognition is the judgment's first exercise, because the test is not always obvious. The senior official's request may be framed as the legitimate government business, and the officer must judge whether the request is the legitimate government business (the public-service product) or the commercial business (the commercial minerals product) in disguise. The recognition's cues include the request's content (the request for the mineral prospectivity is the commercial content, regardless of the framing), the request's beneficiary (the request that benefits a commercial customer is the commercial request, regardless of the requester), and the request's charge code (the request without the commercial charge code is the firewall test, regardless of the official's seniority). The officer must weigh the cues and judge whether the request is the firewall test, and the judgment is the faculty that the experience develops.

The firewall test's resistance is the judgment's second exercise, because the resistance requires the courage to refuse the senior official. The resistance is not the personal refusal but the institutional refusal, with the officer applying the firewall's authority mechanism and referring the request to the office manager. The office manager, who is the firewall's authority, makes the final decision, and the office manager's decision is documented in the firewall log. The officer's resistance is the institutional defence's first line, and the resistance is the condition of the firewall's integrity. An officer who resists the test maintains the firewall; an officer who does not resist erodes the firewall, and the erosion is the failure that the resistance is designed to prevent.

The comparative cases of the firewall being tested are drawn from the published audit reports of the national audit authorities. The United States Government Accountability Office's audits of the United States Geological Survey's commercial activities have identified cases where the firewall was tested by the requests from the senior officials, and the Survey's response, as documented in the audit reports, has been to apply the firewall's authority mechanism and refuse the requests. The National Audit Office's audits of the Ordnance Survey's commercial activities have identified similar cases, and the Ordnance Survey's response has been the same. The cases confirm that the firewall is tested,

and the judgment to recognize and resist the test is a recurring exercise of the geospatial officer's craft, and the exercise is the condition of the firewall's integrity.

34.5 When the Event Demands More Than the Office Can Produce

The fourth situation in which the geospatial officer's judgment is exercised is when a flood event demands a product the office cannot produce in the required time. The situation arises when the event's severity exceeds the office's modeling capacity, when the data acquisition cannot keep pace with the event's evolution, or when the production capacity cannot meet the demand. The officer must assess the office's capacity, determine what can be produced within the required time, and communicate the limitation to the emergency management authority honestly. An officer who overpromises produces intelligence that is late and undermines the authority's trust; an officer who underpromises produces intelligence that is less useful than it could be.

The capacity assessment is the judgment's first exercise, with the officer assessing the office's capacity against the event's demand. The modeling capacity is assessed against the event's severity: the 100-year flood event may be within the modeling capacity, but the 500-year flood event may exceed it; the single-hazard flood event may be within the modeling capacity, but the compound-hazard event (the riverine flooding and the storm surge) may exceed it. The data acquisition capacity is assessed against the event's evolution: the 6-hour update cycle may be achievable with the Sentinel-1 imagery, but the 3-hour update cycle may require the commercial imagery that the office cannot acquire in the required time. The production capacity is assessed against the demand: the office may be able to produce the inundation map and the exposure report, but the office may not be able to produce the evacuation guidance if the evacuation zones' delineation requires the additional analysis. The capacity assessment is the judgment that weighs the office's capabilities against the event's demands, and the assessment is the foundation of the honest communication.

The prioritization is the judgment's second exercise, with the officer prioritizing the products that can be produced within the required time. The prioritization is based on the products' urgency (the inundation map is more urgent than the exposure report, which is more urgent than the evacuation guidance), the products' feasibility (the inundation map is feasible with the available data and the modeling capacity, while the evacuation guidance may not be feasible), and the products' impact (the inundation map has the highest impact on the evacuation orders, while the exposure report has the highest impact on the resource deployment). The officer must weigh the urgency, the feasibility, and the impact, and prioritize the products that maximize the impact within the feasible time. The prioritization is the judgment that allocates the scarce capacity to the most valuable products, and the prioritization is the condition of the office's contribution to the response.

The honest communication is the judgment's third exercise, with the officer communicating the limitation to the emergency management authority. The honest communication comprises the acknowledgment of the limitation (the office cannot produce the evacuation guidance within the required time), the explanation of the limitation's cause (the evacuation zones' delineation requires the additional analysis that the office cannot complete within the time), the offer of the alternative (the office can produce the inundation map and the exposure report, which the authority can use to develop the evacuation zones), and the commitment to the follow-up (the office will produce the evacuation guidance as soon as the additional analysis is completed). The honest communication is

the judgment that maintains the authority's trust, and the communication is the condition of the office's credibility.

The comparative cases of the event demanding more than the office can produce are drawn from the published after-action reviews of the flood events. The Hurricane Katrina after-action reviews (2005) documented the limitations of the geospatial intelligence production during the event, with the Federal Emergency Management Agency's after-action report identifying the gaps in the flood inundation mapping and the levee failure analysis, and the reviews' findings informed the subsequent capacity investments (the USGS's National Flood Hazard Layer, the FEMA's Risk Mapping, Assessment, and Planning program). The European flood events of July 2021 (the floods in Germany, Belgium, and the Netherlands) produced after-action reviews that documented the limitations of the European Flood Awareness System's capacity, with the reviews' findings informing the system's subsequent enhancements (the higher-resolution modeling, the faster update cycles, the improved communication). The cases confirm that the event-demanding-more situation is a recurring situation, and the judgment to assess and communicate the limitation is a recurring exercise of the geospatial officer's craft, and the exercise is the condition of the office's honesty and the authority's trust.

34.6 When the Consumer Requests the Impossible

The fifth situation in which the geospatial officer's judgment is exercised is when the consumer requests the impossible. The situation arises when the consumer (the emergency management authority, the strategic reserve authority, the sectoral regulator, or the water authority) requests a product that the office cannot produce, either because the product is beyond the office's technical capability, or because the product is beyond the available data, or because the product is beyond the time available. The officer must assess the request, determine the office's ability to fulfil the request, and communicate the assessment to the consumer. An officer who accepts the impossible request produces intelligence that is late or wrong; an officer who refuses the impossible request without offering the alternative produces intelligence that is less useful than it could be.

The technical impossibility is the situation where the product is beyond the office's technical capability. The consumer may request the flood inundation map at the 1-metre resolution, but the office's hydraulic model is at the 2-metre resolution, and the 1-metre resolution is beyond the model's capability. The consumer may request the crop yield forecast at the parcel level, but the office's crop model is at the district level, and the parcel level is beyond the model's capability. The consumer may request the mineral prospectivity at the 1:10,000 scale, but the office's prospectivity model is at the 1:100,000 scale, and the 1:10,000 scale is beyond the model's capability. The officer must assess the technical impossibility, communicate the limitation, and offer the alternative (the product at the office's capability, with the explanation of the capability's limit).

The data impossibility is the situation where the product is beyond the available data. The consumer may request the flood inundation map for the area that is covered by the cloud, but the optical imagery cannot penetrate the cloud, and the radar imagery may not be available for the required time. The consumer may request the crop condition assessment for the area that is not covered by the Sentinel-2 imagery, but the Sentinel-2's revisit time may not provide the coverage for the required period. The consumer may request the change detection for the area that is not covered by the high-resolution imagery, but the commercial imagery may not be available for the required

area and the required time. The officer must assess the data impossibility, communicate the limitation, and offer the alternative (the product with the available data, with the explanation of the data's limit).

The time impossibility is the situation where the product is beyond the time available. The consumer may request the flood inundation map within 1 hour of the activation, but the office's modeling and the production require 4 hours, and the 1-hour delivery is beyond the office's capability. The consumer may request the crop yield forecast within 1 day of the harvest, but the office's crop model and the validation require 1 week, and the 1-day delivery is beyond the office's capability. The consumer may request the change detection within 1 day of the satellite acquisition, but the office's processing and the analysis require 3 days, and the 1-day delivery is beyond the office's capability. The officer must assess the time impossibility, communicate the limitation, and offer the alternative (the product within the office's capability, with the explanation of the time's limit).

The alternative is the judgment's substantive output, with the officer offering the consumer the product that the office can produce, rather than the product that the consumer requested. The alternative must be relevant to the consumer's need (the product that addresses the consumer's underlying question, even if the product is not the requested product), feasible for the office (the product that the office can produce within the technical, the data, and the time constraints), and clearly communicated (the explanation of the alternative's relationship to the requested product, the alternative's limitations, and the alternative's value). The alternative is the judgment that maintains the consumer's trust while being honest about the office's limitations, and the alternative is the condition of the office's contribution to the consumer's decisions.

The comparative cases of the consumer requesting the impossible are drawn from the published after-action reviews and the professional literature. The Copernicus Emergency Management Service's after-action reviews have documented the cases where the service's activation requests were beyond the service's capability (the requests for the 1-metre resolution damage assessment, which the service could not provide with the Sentinel imagery), and the service's response was to offer the 10-metre resolution assessment with the explanation of the resolution's limit. The United States Geological Survey's after-action reviews have documented the cases where the survey's product requests were beyond the survey's capability (the requests for the real-time flood inundation mapping, which the survey could not provide with the available modeling and the data), and the survey's response was to offer the near-real-time mapping with the explanation of the time's limit. The cases confirm that the consumer requesting the impossible is a recurring situation, and the judgment to assess and offer the alternative is a recurring exercise of the geospatial officer's craft.

34.7 The Development of Judgment

The development of the geospatial officer's judgment is the process by which the officer acquires the craft, and the process is the combination of the experience, the study, and the mentorship. The experience is the accumulation of the situations in which the judgment is exercised, and the experience is the foundation of the judgment. The study is the examination of the comparative cases, through the after-action reviews, the professional literature, and the lessons-learned documents, and the study is the mechanism by which the officer learns from the experience of others. The mentorship is the guidance of the senior officers, who share their judgment with the junior officers through the

working relationships, and the mentorship is the mechanism by which the judgment is transmitted across the generations.

The experience is the judgment's primary source, with the officer encountering the situations, exercising the judgment, and learning from the encounter. The experience is accumulated over the career, with the junior officers encountering the simpler situations (the data ambiguity, the model disagreement) and the senior officers encountering the more complex situations (the firewall test, the event demand, the consumer request). The experience's accumulation is the reason the judgment is the faculty of the experienced officer, and the experience is the reason the judgment cannot be taught in the classroom. The office that provides its officers with the diverse experiences (the rotations across the domains, the assignments to the different roles, the deployments to the events) accelerates the judgment's development, and the diverse experiences are the condition of the judgment's breadth.

The study is the judgment's secondary source, with the officer learning from the experience of others through the comparative cases. The study comprises the after-action reviews (the reviews of the events, the incidents, and the operations, treated in Chapter 32, section 32.4), the professional literature (the journals, the conference proceedings, the technical reports), and the lessons-learned documents (the internal documents that capture the office's lessons). The study is the mechanism by which the officer learns from the comparative cases without personally encountering the situations, and the study is the mechanism by which the officer's judgment is informed by the broader profession's experience. The office that encourages its officers' study (the time for the reading, the budget for the conference attendance, the support for the publication) accelerates the judgment's development, and the study is the condition of the judgment's depth.

The mentorship is the judgment's tertiary source, with the senior officers sharing their judgment with the junior officers through the working relationships. The mentorship comprises the formal mentorship (the assigned mentor-mentee relationships, with the regular meetings and the structured guidance) and the informal mentorship (the daily working relationships, with the senior officers providing the guidance in the course of the work). The mentorship is the mechanism by which the judgment is transmitted across the generations, and the mentorship is the reason the judgment is the profession's collective faculty rather than the individual officer's personal faculty. The office that fosters the mentorship (the formal mentorship programme, the informal mentorship culture, the senior officers' engagement with the junior officers) accelerates the judgment's development, and the mentorship is the condition of the judgment's transmission.

The development of the judgment is not a matter of the training courses or the certifications, which cover the rules but not the judgment. The development is a matter of the career, and the career is the process by which the officer encounters the situations, exercises the judgment, and learns from the encounter. The office that develops its officers' judgment produces officers who can navigate the situations where the rules are silent. The office that does not develop the judgment produces officers who are confined to the rules, with the confinement limiting the officers' ability to address the situations the rules do not cover. The development of the judgment is the office's investment in its officers' craft, and the investment is the condition of the office's resilience.

The judgment's development is assessed through the performance review, with the office manager evaluating the officer's judgment in the situations the officer has encountered. The performance review's judgment assessment is based on the decision's quality (the decision was

correct, given the information available), the decision's timeliness (the decision was made in the time required), the decision's documentation (the decision was documented for the review and the audit), and the decision's communication (the decision was communicated to the affected parties). The judgment assessment is the input to the officer's development plan, with the plan identifying the experiences, the study, and the mentorship that will develop the officer's judgment further. The judgment assessment is the mechanism by which the office tracks the judgment's development, and the assessment is the condition of the office's systematic investment in the judgment.

34.8 The Judgment and the Rules

The judgment and the rules are not in opposition; they are complementary. The rules cover the routine work, and the judgment covers the situations where the rules are silent. An officer who relies only on the rules is confined to the routine; an officer who relies only on the judgment is unmoored from the standards. An officer who combines the rules and the judgment produces work that is both standard and nuanced, and the combination is what distinguishes the professional from the technician. The rules are the foundation, and the judgment is the edifice, and the two together are the geospatial officer's craft.

The rules are the codified experience of the profession, with the rules developed through the comparative practice (treated in Chapter 33), the lessons learned (treated in Chapter 32), and the professional standards. The rules comprise the doctrine (the doctrinal principles treated throughout this Handbook), the procedures (the operational procedures treated in Part III and the inter-agency procedures treated in Part V), and the job breakdown sheets (the step-by-step instructions treated in Part III). The rules are the profession's collective judgment, codified for the routine application, and the rules are the mechanism by which the profession's experience is transmitted to the new officers.

The judgment is the individual officer's faculty, with the judgment developed through the experience, the study, and the mentorship (treated in section 34.7). The judgment comprises the diagnostic faculty (the ability to diagnose the situation's cause), the evaluative faculty (the ability to evaluate the options' merits), the decisive faculty (the ability to decide on the option), and the communicative faculty (the ability to communicate the decision). The judgment is the individual officer's contribution to the situations the rules do not cover, and the judgment is the mechanism by which the officer navigates the ambiguity, the consequence, and the accountability that the rules do not resolve.

The relationship between the judgment and the rules is the relationship that the Handbook is designed to support. The Handbook provides the rules, through the doctrine, the procedures, and the job breakdown sheets, and the Handbook provides the judgment, through the comparative cases, the after-action reviews, and the professional literature. An officer who studies the Handbook acquires both the rules and the judgment, and the acquisition is the condition of the officer's competence. The judgment is not a substitute for the rules, and the rules are not a substitute for the judgment, and an officer who treats either as a substitute for the other has misunderstood the craft.

The rules and the judgment evolve together, with the rules expanding to cover the situations that were previously the judgment's domain, and the judgment expanding to address the new situations that the rules do not yet cover. The rules' expansion is the process by which the profession codifies its experience, with the rules developed for the situations that the officers have encountered and the

judgment has resolved. The judgment's expansion is the process by which the profession encounters the new situations, with the judgment exercised for the situations that the rules do not yet cover. The two expansions together constitute the profession's evolution, and the evolution is the mechanism by which the profession advances.

34.9 Worked Example: Exercising Judgment During the Riverbourne Basin Flood Event

The judgment's exercise is illustrated through the Riverbourne Basin flood event of March 2024, treated in Chapters 24, 27, 28, 29, 32, and 33. The event's response required the officers' judgment in the multiple situations, with each situation testing a different judgment faculty. The worked example traces the event's timeline through the five judgment situations treated in sections 34.2 through 34.6, demonstrating how the officers exercised the judgment and how the judgment's exercise contributed to the event's response.

The first judgment situation (the model-ground truth disagreement, treated in section 34.2) is encountered at 06:00 on 25 March 2024, when the hazard specialist Dr. Maren Holloway (treated in Chapter 30, section 30.9) compares the HEC-RAS model's predicted peak stage with the observed peak stage. The model predicted the peak stage of 4.05 metres above the flood stage, and the observed peak stage (from the Portsmith gauge) is 4.2 metres, a difference of 0.15 metres. Additionally, the model predicted the no-breach inundation extent of 12 square kilometres, and the observed inundation extent (from the Sentinel-1 radar imagery acquired at 04:30) is 18 square kilometres, an underestimation of 33 per cent. The hazard specialist must diagnose the cause of the disagreement: is it a model error (the model's parameters are wrong), a data error (the input data is wrong), a validation error (the observed data is wrong), or a genuine change (the conditions have changed)?

The hazard specialist's diagnosis begins with the data error check: the rainfall data is verified against the rain gauge observations (the rain gauge data confirms the radar rainfall estimates, ruling out the rainfall data error), the river stage data is verified against the gauge's datum (the datum is correct, ruling out the river stage data error), and the soil moisture data is verified against the in-situ sensors (the sensors confirm the satellite data, ruling out the soil moisture data error). The data error is ruled out. The validation error check follows: the observed peak stage is verified against the gauge's record (the gauge was not overtopped, ruling out the gauge error), and the observed inundation extent is verified against the field reports (the field reports confirm the satellite observation, ruling out the satellite error). The validation error is ruled out. The model error check follows: the HEC-RAS model's parameters are reviewed (the Manning's n values, the levee data, the channel geometry are all correct, ruling out the model error). The model error is ruled out. The genuine change is the remaining cause, and the hazard specialist identifies the genuine change: the storm surge's contribution was larger than the forecast, because the deterministic forecast underestimated the storm surge by 20 per cent. The genuine change's response is the model's update, with the storm surge boundary condition updated to use the observed storm surge data, and the revised model run produces the revised inundation map at 06:30 on 25 March. The judgment's exercise produced the correct diagnosis (the genuine change) and the correct response (the model's update), and the exercise is the illustration of the diagnostic faculty.

The second judgment situation (the data ambiguity, treated in section 34.3) is encountered at 04:30 on 25 March 2024, when the data pipeline operator acquires the Sentinel-1 radar imagery for the inundation assessment. The radar imagery is the only available imagery, because the optical imagery (the Sentinel-2) is blocked by the cloud cover that accompanies the flood event. The radar imagery, however, has the artifacts: the layover from the buildings in the Portsmith urban area, the shadow from the hills in the Riverbourne valley, and the speckle from the coherent illumination. The hazard specialist must judge whether the artifacts are the real features (the inundation) or the artifacts (the false positives), and the judgment requires the understanding of the radar imaging geometry and the radar processing.

The hazard specialist's judgment begins with the artifact identification: the layover is identified by the geometric relationship between the sensor's look direction and the buildings (the layover appears on the side of the buildings facing the sensor, and the inundation appears in the flat areas), the shadow is identified by the geometric relationship between the sensor's look direction and the hills (the shadow appears behind the hills, and the inundation appears in the valleys), and the speckle is identified by the granular pattern (the speckle is the random pattern, and the inundation is the coherent pattern). The artifacts are distinguished from the inundation, and the inundation extent is mapped from the radar imagery with the artifacts excluded. The judgment's exercise produced the correct distinction (the artifacts vs the inundation) and the correct mapping (the 18-square-kilometre inundation extent), and the exercise is the illustration of the evaluative faculty.

The third judgment situation (the firewall test, treated in section 34.4) is encountered at 14:00 on 24 March 2024, when a senior official from the Ministry of Energy calls the office manager and requests the "briefing" on the mineral prospectivity of the Westbourne area, arguing that the briefing is for the government's energy security assessment and should be provided under the public-service appropriation. The office manager must judge whether the request is the legitimate government business (the public-service product) or the commercial business (the commercial minerals product) in disguise, and the judgment requires the recognition of the firewall test's cues.

The office manager's judgment begins with the request's content: the mineral prospectivity is the commercial content (treated in Chapter 22 and Chapter 26), regardless of the framing as the "briefing". The request's beneficiary is the Ministry of Energy, but the Ministry of Energy's energy security assessment may benefit the commercial minerals customers (the Westbourne Minerals Company, treated in Chapter 32, section 32.9, which holds the exploration licence in the Westbourne area). The request's charge code is absent, which is the firewall test's trigger. The office manager recognizes the firewall test, applies the firewall's authority mechanism (the charge code verification, the office manager's authorization, the documentation), and communicates to the senior official that the mineral prospectivity is the commercial product, and the commercial charge code is required for the release. The senior official accepts the response, and the firewall test is documented in the office's firewall log. The judgment's exercise produced the correct recognition (the firewall test) and the correct response (the firewall's authority mechanism), and the exercise is the illustration of the decisive faculty.

The fourth judgment situation (the event demanding more than the office can produce, treated in section 34.5) is encountered at 10:00 on 24 March 2024, when the emergency management authority's emergency operations centre requests the evacuation guidance for the 6 evacuation zones

within 2 hours. The office manager must assess the office's capacity to produce the evacuation guidance within the required time, and the assessment requires the judgment to weigh the office's capabilities against the authority's demands.

The office manager's capacity assessment begins with the modeling capacity: the evacuation guidance requires the delineation of the evacuation zones, which requires the inundation map's update and the evacuation route's assessment, and the update and the assessment require 4 hours (the HEC-RAS model's update takes 90 minutes, the evacuation route's assessment takes 60 minutes, and the production takes 90 minutes). The 2-hour delivery is beyond the office's capacity. The data acquisition capacity is assessed: the Sentinel-1 imagery for the update is available, but the high-resolution imagery for the evacuation route's assessment is not available within the 2-hour window. The production capacity is assessed: the office has the hazard specialist and the production operator on duty, but the evacuation guidance's production requires the additional analysis that the office cannot complete within the 2-hour window. The capacity assessment concludes that the evacuation guidance cannot be produced within the 2-hour window.

The office manager's prioritization follows: the inundation map (already produced and delivered at 12:30 on 23 March) and the exposure report (already produced and delivered at 13:00 on 23 March) are the products that the authority can use to develop the evacuation zones, and the evacuation guidance can be produced within the 4-hour window. The office manager communicates the limitation to the authority's emergency operations centre at 10:30 on 24 March: the acknowledgment (the evacuation guidance cannot be produced within the 2-hour window), the explanation (the evacuation zones' delineation requires the additional analysis that the office cannot complete within the time), the offer of the alternative (the inundation map and the exposure report are available, which the authority can use to develop the evacuation zones), and the commitment to the follow-up (the evacuation guidance will be produced within the 4-hour window, by 14:00 on 24 March). The authority's emergency operations centre accepts the response, and the evacuation guidance is delivered at 14:00 on 24 March as committed. The judgment's exercise produced the correct assessment (the capacity limitation), the correct prioritization (the alternative products), and the correct communication (the honest acknowledgment with the alternative), and the exercise is the illustration of the communicative faculty.

The fifth judgment situation (the consumer requesting the impossible, treated in section 34.6) is encountered at 16:00 on 25 March 2024, when the water authority's operations director requests the levee breach probability assessment for the Levee Section 23 (treated in Chapter 27, section 27.9) within 1 hour. The office manager must assess the office's ability to fulfil the request, and the assessment requires the judgment to determine the request's feasibility.

The office manager's assessment begins with the technical feasibility: the levee breach probability assessment requires the geotechnical analysis of the levee's foundation and the seepage's progression, and the geotechnical analysis is beyond the office's technical capability (the office's hazard specialist is the hydrologist, not the geotechnical engineer). The data feasibility is assessed: the levee's geotechnical data (the foundation's soil properties, the seepage's flow rate) is not available to the office (the data is held by the water authority's geotechnical team). The time feasibility is assessed: the 1-hour delivery is beyond the office's capability, even if the technical and the data

feasibility were available. The assessment concludes that the request is the impossible request (the technical, the data, and the time impossibility).

The office manager's alternative follows: the office can offer the levee breach inundation map (the map of the inundation that would result from the Levee Section 23's breach, which the office has already produced as the breach scenario at 12:30 on 23 March, treated in Chapter 27, section 27.9), which addresses the water authority's underlying question (the consequence of the breach) even if the product is not the requested product (the breach probability). The office manager communicates the assessment to the water authority's operations director at 16:30 on 25 March: the acknowledgment (the breach probability assessment cannot be produced within the 1-hour window), the explanation (the geotechnical analysis is beyond the office's technical capability, and the geotechnical data is not available), the offer of the alternative (the breach inundation map is available, which shows the consequence of the breach), and the recommendation (the water authority's geotechnical team should conduct the breach probability assessment, with the office's breach inundation map as the input). The water authority's operations director accepts the response, and the geotechnical team conducts the breach probability assessment using the office's breach inundation map as the input. The judgment's exercise produced the correct assessment (the impossible request), the correct alternative (the breach inundation map), and the correct communication (the honest acknowledgment with the alternative), and the exercise is the illustration of the communicative faculty.

This worked example illustrates the judgment's full exercise: the diagnostic faculty (the model-ground truth disagreement, the genuine change diagnosis, the model's update), the evaluative faculty (the data ambiguity, the artifact distinction, the inundation mapping), the decisive faculty (the firewall test, the recognition and the resistance, the authority mechanism's application), the communicative faculty (the event demand, the capacity assessment, the honest communication with the alternative), and the communicative faculty (the consumer request, the impossible request assessment, the alternative offer). The five judgment situations were encountered across the event's timeline (the model-ground truth disagreement at 06:00 on 25 March, the data ambiguity at 04:30 on 25 March, the firewall test at 14:00 on 24 March, the event demand at 10:00 on 24 March, the consumer request at 16:00 on 25 March), and each situation tested a different judgment faculty. The judgment's exercise contributed to the event's response: the model's update produced the revised inundation map, the artifact distinction produced the accurate inundation extent, the firewall's resistance maintained the commercial firewall's integrity, the honest communication maintained the authority's trust, and the alternative offer maintained the water authority's trust. The judgment is the craft dimension of the geospatial profession, and the craft is the dimension that makes the profession a profession rather than a trade. The officers who exercised the judgment (Dr. Maren Holloway and the office manager) are the experienced officers who have developed the judgment through the experience, the study, and the mentorship, and the development is the reason the judgment was available when the situations arose. The judgment is the faculty that the Handbook is designed to develop, and the development is the condition of the office's ability to navigate the situations where the rules are silent, and the navigation is the condition of the office's contribution to the public safety during the flood events.

Chapter 35 — Drafting and Communication for Geospatial Intelligence

35.1 The Variation by Domain

Drafting and communication for geospatial intelligence varies by domain and by the public-commercial boundary. The register of a flood bulletin differs from that of a procurement advisory, which differs from that of an energy infrastructure risk assessment, which differs from that of a commercial prospectivity report. The variation is not stylistic preference but functional necessity: each product type serves a different purpose, addresses a different audience, and operates under a different classification regime. An officer who drafts all products in the same register produces products that are inappropriate for their purposes, with the inappropriateness undermining the products' usefulness and the consumers' trust.

The flood bulletin is drafted in the operational register, because the bulletin's purpose is to support the emergency management authority's operational response, and the authority needs the intelligence in a form that is concise, clear, and actionable. The procurement advisory is drafted in the analytical register, because the advisory's purpose is to inform the strategic reserve authority's decision, and the authority needs the intelligence in a form that is thorough, balanced, and honest about the uncertainties. The energy infrastructure risk assessment is drafted in the security register, because the assessment's purpose is to support the energy sector entity's infrastructure protection, and the entity needs the intelligence in a form that is precise, classified, and actionable. The commercial prospectivity report is drafted in the commercial register, because the report's purpose is to serve the commercial customer's exploration decision, and the customer needs the intelligence in a form that is professional, confident, and compliant with the reporting standards.

The four registers are distinguished by four characteristics. The first characteristic is the audience: the operational register addresses the emergency responders, the analytical register addresses the decision-makers, the security register addresses the security-cleared personnel, and the commercial register addresses the commercial customers. The second characteristic is the purpose: the operational register supports the response, the analytical register informs the decision, the security register protects the infrastructure, and the commercial register serves the customer. The third characteristic is the classification: the operational register ranges from Public to Internal, the analytical register ranges from Public to Internal, the security register is Security Confidential, and the commercial register is Confidential. The fourth characteristic is the timeline: the operational register is event-driven (hours), the analytical register is seasonal (weeks), the security register is monthly, and the commercial register is project-based (weeks to months). The four characteristics together define each register, and the characteristics are the conditions that the drafting must respect.

The drafting across the four registers is the craft that this chapter develops. The craft comprises the register selection (the choice of the appropriate register for the product), the register execution (the drafting in the selected register), and the register verification (the review that verifies the register's appropriateness). The craft is developed through the experience (the drafting of the multiple products across the registers), the study (the examination of the comparative products), and the mentorship (the guidance of the senior officers). The craft is the communication dimension of the

geospatial profession, and the craft is the dimension that the training (treated in Chapter 31) and the judgment (treated in Chapter 34) support.

35.2 The Flood Bulletin

The flood bulletin is the office's event-driven product, and its drafting is governed by the operational register. The bulletin comprises the event summary, the inundation map, the exposure assessment, and the evacuation guidance. The event summary states the event's nature, severity, and expected duration in concise, declarative sentences. The inundation map shows the flooded area, the depth, and the extent in the standard cartographic format. The exposure assessment quantifies the population, property, and infrastructure affected, in the operational format that the emergency management authority's systems can ingest. The evacuation guidance identifies the evacuation routes and the shelter locations, in the public safety format that the authority can communicate to the public.

The event summary is the bulletin's first section, and it is the section that the authority's decision-makers read first. The summary comprises the event's nature (the type of the flood: the riverine flooding, the coastal flooding, the compound flooding), the event's severity (the peak stage, the return period, the comparison with the historical events), and the event's expected duration (the time to the peak, the time to the recession, the time to the stand-down). The summary is drafted in the concise, declarative sentences, with each sentence conveying one piece of the information. The summary's example: "The Riverbourne at Portsmith is expected to peak at 4.2 metres above the flood stage at 04:00 on 25 March 2024, exceeding the previous record of 3.52 metres set in 2019. The event is a compound flood, combining the riverine flooding from the 72-hour rainfall of 215 millimetres and the coastal flooding from the storm surge of 2.8 metres. The event is expected to recede by 28 March 2024." The summary's concision is the condition of the authority's rapid comprehension.

The inundation map is the bulletin's second section, and it is the section that the authority's operations officers use for the evacuation orders and the resource deployment. The map shows the flooded area (the inundation extent, colour-coded by the depth: 0-0.5 metres, 0.5-1.0 metres, 1.0-2.0 metres, 2.0 metres and above), the critical infrastructure (the hospitals, the substations, the water-treatment plants, the hazardous-materials sites), and the evacuation zones (the recommended evacuation areas, colour-coded by the priority). The map is produced in the standard cartographic format (the GeoTIFF for the authority's GIS, the PDF for the decision-makers, the KML for the field teams), with the Internal classification marking (and the Security Confidential marking for the version that includes the critical infrastructure).

The exposure assessment is the bulletin's third section, and it is the section that the authority's planning officers use for the resource allocation and the public communication. The assessment quantifies the population exposed (the total number, the breakdown by the age and the vulnerability), the property exposed (the number of the buildings, the breakdown by the type: the residential, the commercial, the industrial, the public), and the infrastructure exposed (the number and the type of the critical infrastructure, the hazardous-materials sites). The assessment is produced in the operational format (the CSV for the authority's systems, the PDF for the decision-makers), with the Internal classification marking.

The evacuation guidance is the bulletin's fourth section, and it is the section that the authority's emergency operations centre uses for the evacuation orders. The guidance identifies the evacuation

zones (the 6 zones, with the clear geographical boundaries: the roads, the rivers, the administrative boundaries), the evacuation routes (the roads that remain passable, with the capacity and the estimated travel time), and the shelter locations (the designated shelters, with the capacity and the current occupancy). The guidance is produced in the public safety format (the PDF for the authority, the KML for the field teams, the Public version for the public communication).

The flood bulletin is drafted under time pressure, and the drafting discipline is the discipline of producing the bulletin that is clear, accurate, and complete within the time available. The officer who drafts the bulletin must resist the pressure to sacrifice the clarity for the speed, because the bulletin that is fast but unclear is not useful. The officer must resist the pressure to sacrifice the accuracy for the completeness, because the bulletin that is complete but inaccurate is misleading. The drafting discipline is the discipline of the operational register, and the discipline is the condition of the bulletin's usefulness for the emergency response.

The flood bulletin's drafting is reviewed by the office manager before the release, with the review verifying the bulletin's accuracy (the data is correct), the bulletin's clarity (the language is concise and declarative), the bulletin's completeness (the four sections are present and the required information is included), and the bulletin's classification (the marking is applied correctly). The review is conducted within the 30-minute window before the release, with the review's efficiency supported by the bulletin's template (the standard structure that the officer fills in) and the office manager's familiarity with the bulletin's format. The review is the final check that ensures the bulletin's quality, and the review is the condition of the bulletin's credibility.

35.3 The Procurement Advisory

The procurement advisory is the office's forward-sensor product, and its drafting is governed by the analytical register. The advisory comprises the yield forecast, the stock level assessment, the supply-demand analysis, and the release or buy signal. The yield forecast states the forecast and its assumptions and uncertainties. The stock level assessment states the current stock levels and their source. The supply-demand analysis synthesizes the forecast and the stock levels with the demand assumptions. The release or buy signal states the advisory's conclusion in clear, direct language.

The yield forecast section is the advisory's first section, and it is the section that the reserve authority's procurement director reads to understand the season's crop outlook. The section comprises the forecast (the yield in tonnes per hectare, by district and by region, with the confidence interval), the assumptions (the weather scenario, the soil moisture, the management conditions), and the uncertainties (the confidence interval's basis, the forecast's sensitivity to the weather, the forecast's comparison with the previous forecast). The section is drafted in the analytical register, with the thorough presentation of the forecast and the honest presentation of the uncertainties. The section's example: "The Northfields' wheat yield is forecast at 2.1 tonnes per hectare, with the 90-per-cent confidence interval of 1.8 to 2.4 tonnes per hectare. The forecast is based on the WOFOST crop model, configured with the observed rainfall (65 per cent of normal for March-July 2024), the soil moisture (25 per cent below normal from the NASA SMAP satellite), and the weather outlook (the below-normal rainfall and the above-normal temperatures expected to continue). The forecast is a 19-per-cent downward revision from the planting forecast of 2.6 tonnes per hectare, reflecting the drought's intensification."

The stock level assessment section is the advisory's second section, and it is the section that the procurement director reads to understand the reserve's current position. The section comprises the current stock levels (by commodity, by storage facility, by quality grade), the stock-to-consumption ratio (the months of consumption that the stocks represent), and the stock levels' trend (the comparison with the previous months and the historical patterns). The section is drafted in the analytical register, with the thorough presentation of the stock levels and the honest presentation of the trend. The section's example: "The reserve's wheat stocks are 1.85 million tonnes, distributed across the four storage depots (the Northfields Depot 0.65, the Portsmith Depot 0.50, the Capital Depot 0.40, the Western Depot 0.30). The stock-to-consumption ratio is 3.7 months, below the authority's target of 4.5 months. The stock levels have declined by 12 per cent over the past 12 months, reflecting the higher-than-expected consumption and the lower-than-expected procurement."

The supply-demand analysis section is the advisory's third section, and it is the section that the procurement director reads to understand the season's balance. The section comprises the production estimate (the yield multiplied by the planted area), the consumption estimate (the population multiplied by the per-capita consumption rate), the annual balance (the production plus the carry-over minus the consumption), and the reserve's projected level (the current stocks plus the procurement minus the consumption until the harvest). The section is drafted in the analytical register, with the thorough presentation of the analysis and the honest presentation of the balance's implications.

The release or buy signal section is the advisory's fourth section, and it is the section that the procurement director reads to understand the advisory's recommendation. The section comprises the signal (the release, the buy, or the hold), the recommended quantity (the tonnes to release or buy), the rationale (the synthesis of the yield forecast, the stock levels, and the supply-demand analysis), and the timing (the recommended execution window). The section is drafted in the analytical register, with the direct presentation of the signal and the thorough presentation of the rationale. The section's example: "The advisory recommends a buy of 0.20 million tonnes, to be executed before the harvest (by 31 August 2024). The buy is recommended because the drought-reduced yield (2.1 tonnes per hectare, 19 per cent below the planting forecast) and the below-target stock-to-consumption ratio (3.7 months, below the 4.5-month target) warrant the precautionary procurement to maintain the reserve's adequacy through the next consumption cycle."

The procurement advisory is drafted in the analytical register, and the drafting discipline is the discipline of producing the advisory that is thorough, balanced, and honest about the uncertainties. The officer who drafts the advisory must resist the pressure to state the forecast with false precision, because the forecast that is precise but wrong is misleading. The officer must resist the pressure to state the signal without the supporting analysis, because the signal without the analysis is not an advisory but a directive. The drafting discipline is the discipline of the analytical register, and the discipline is the condition of the advisory's usefulness for the reserve authority's decision.

The procurement advisory's drafting is reviewed by the office manager before the release, with the review verifying the advisory's accuracy (the forecast and the stock levels are correct), the advisory's thoroughness (the four sections are present and the analysis is complete), the advisory's balance (the uncertainties are honestly presented), and the advisory's classification (the Internal marking is applied). The review is conducted within the 2-hour window before the release, with the

review's efficiency supported by the advisory's template. The review is the final check that ensures the advisory's quality, and the review is the condition of the advisory's credibility.

35.4 The Energy Infrastructure Risk Assessment

The energy infrastructure risk assessment is the office's security-confidential product, and its drafting is governed by the security register. The assessment comprises the infrastructure risk map, the encroachment alerts, and the maintenance advisories. The infrastructure risk map shows the monitored infrastructure and the vulnerability scores in the standard cartographic format, with the Security Confidential marking. The encroachment alerts identify the detected encroachments and their threat to the infrastructure. The maintenance advisories identify the infrastructure at risk from the ground conditions and the recommended maintenance.

The infrastructure risk map section is the assessment's first section, and it is the section that the energy sector entity's operations centre and the sectoral regulator's compliance division use for the infrastructure management and the compliance verification. The section comprises the monitored infrastructure (the generation facilities, the transmission corridors, the substations, the pipelines, overlaid on the high-resolution satellite imagery), the vulnerability scores (colour-coded by priority: red for high-priority, orange for moderate-priority, yellow for low-priority), and the assessment's date and the coverage. The section is drafted in the security register, with the precise identification of the infrastructure's location and the precise assessment of the vulnerability. The section's example: "The Highmoor-Portsmith transmission corridor (the 220-kilometre, 500-kV transmission line from the Highmoor Power Station to the Portsmith Substation) is monitored as of 5 July 2024. The corridor's vulnerability scores are: 4 high-priority (the vegetation encroachments within the 15-metre clearance zone, near the critical towers), 3 moderate-priority (the third-party constructions within the 30-metre right-of-way), and 1 high-priority (the ground subsidence near the critical tower)."

The encroachment alerts section is the assessment's second section, and it is the section that the entity's maintenance dispatch uses for the encroachment response. The section comprises the alert's identifier, the encroachment's type (the vegetation encroachment, the third-party construction, the ground subsidence, the excavation), the encroachment's location (the latitude, the longitude, the nearest tower identifier, the distance to the tower), the encroachment's detection date and method, the vulnerability score, and the recommended action. The section is drafted in the security register, with the precise identification of each encroachment and the clear recommendation for the response.

The maintenance advisories section is the assessment's third section, and it is the section that the entity's maintenance planning uses for the longer-term maintenance. The section comprises the advisory's identifier, the infrastructure at risk (the section of the corridor, the specific towers or the structures), the risk's cause (the ground subsidence, the erosion, the slope instability), the recommended maintenance (the vegetation clearing, the ground stabilization, the structural reinforcement), and the recommended timeline. The section is drafted in the security register, with the precise identification of the infrastructure at risk and the clear recommendation for the maintenance.

The energy infrastructure risk assessment is drafted in the security register, and the drafting discipline is the discipline of producing the assessment that is precise, classified, and actionable. The officer who drafts the assessment must ensure that the classification marking is applied to every

component, because the assessment that is unmarked or inconsistently marked is a security breach. The officer must ensure that the assessment is precise about the infrastructure's location and vulnerability, because the assessment that is imprecise is not actionable. The drafting discipline is the discipline of the security register, and the discipline is the condition of the assessment's usefulness for the infrastructure's protection.

The assessment's drafting is reviewed by the office manager before the release, with the review verifying the assessment's accuracy (the encroachments and the vulnerability scores are correct), the assessment's precision (the infrastructure's location and the vulnerability are precisely identified), the assessment's classification (the Security Confidential marking is applied to every component), and the assessment's actionability (the recommended actions are clear and feasible). The review is conducted within the 1-day window before the release, with the review's thoroughness reflecting the assessment's security sensitivity. The review is the final check that ensures the assessment's quality and the security, and the review is the condition of the assessment's credibility and the infrastructure's protection.

35.5 The Commercial Prospectivity Report

The commercial prospectivity report is the office's commercial product, and its drafting is governed by the commercial register. The report comprises the prospectivity map, the resource estimate, and the geological interpretation. The prospectivity map shows the mineral potential in the standard cartographic format, with the Confidential marking. The resource estimate states the estimated quantity and quality of the mineral resource, in the format that complies with the international reporting standards. The geological interpretation explains the prospectivity assessment's geological basis, in the professional language that the commercial customer's geologists can evaluate.

The prospectivity map section is the report's first section, and it is the section that the customer's exploration team uses for the target selection. The section comprises the prospectivity model (the knowledge-driven or the data-driven model, with the weights and the evidence layers), the prospectivity map (the mineral potential, colour-coded by the prospectivity class: the high, the moderate, the low), the target areas (the specific areas recommended for the follow-up exploration), and the map's date and the coverage. The section is drafted in the commercial register, with the professional presentation of the prospectivity and the confident assessment of the potential. The section's example: "The Westbourne area's copper prospectivity is assessed using the weights-of-evidence model, with the evidence layers comprising the geological maps, the geochemical data, the geophysical data, and the known mineral occurrences. The prospectivity map identifies 3 high-prospectivity targets (the Westbourne North, the Westbourne Central, and the Westbourne South), with the combined area of 15 square kilometres, recommended for the follow-up exploration."

The resource estimate section is the report's second section, and it is the section that the customer's resource geologists use for the resource declaration and the mining investment. The section comprises the resource estimate (the tonnage and the grade, by the resource class: the measured, the indicated, the inferred), the estimation method (the kriging or the inverse distance weighting, with the parameters), the resource's confidence (the geological confidence and the estimation confidence), and the compliance statement (the compliance with the international

reporting standards). The section is drafted in the commercial register, with the professional presentation of the resource and the confident assessment of the estimate. The section must comply with the international reporting standards: the Committee for Mineral Reserves International Reporting Standards (CRIRSCO) framework, the Australasian Code for Reporting of Exploration Results, Mineral Resources and Ore Reserves (the JORC Code, 2012 edition), the Canadian Institute of Mining, Metallurgy and Petroleum Definition Standards (the CIM Standards, 2014 edition), or the South African Code for the Reporting of Exploration Results, Mineral Resources and Mineral Reserves (the SAMREC Code, 2016 edition).

The geological interpretation section is the report's third section, and it is the section that the customer's geologists use for the geological evaluation. The section comprises the regional geology (the geological setting, the stratigraphy, the structural geology), the local geology (the specific geological features of the target areas), the mineralization (the mineralization's type, the mineralization's controls, the mineralization's timing), and the exploration potential (the potential for the additional discoveries, the recommended exploration activities). The section is drafted in the commercial register, with the professional presentation of the geology and the confident assessment of the potential. The section's language is the professional geological language, with the technical terms (the lithology, the alteration, the structural control, the paragenesis) used precisely and the geological reasoning presented clearly.

The commercial prospectivity report is drafted in the commercial register, and the drafting discipline is the discipline of producing the report that is professional, confident, and compliant with the reporting standards. The officer who drafts the report must ensure that the report is professional in its presentation, because the report that is unprofessional undermines the commercial arm's credibility. The officer must ensure that the report is confident in its assessment, because the report that is hesitant is not useful for the customer's exploration decision. The officer must ensure that the report is compliant with the international reporting standards, because the report that is non-compliant exposes the office to legal liability. The drafting discipline is the discipline of the commercial register, and the discipline is the condition of the report's commercial viability.

The report's drafting is reviewed by the office manager before the release, with the review verifying the report's professionalism (the presentation is professional), the report's confidence (the assessment is confident), the report's compliance (the reporting standards are followed), and the report's classification (the Confidential marking is applied). The review is conducted within the 1-week window before the release, with the review's thoroughness reflecting the report's commercial value. The review is the final check that ensures the report's quality and the commercial viability, and the review is the condition of the report's credibility and the commercial arm's sustainability.

35.6 The Government Style Manual and the Plain Language Standard

The government style manual is the authoritative reference for the drafting of government documents, and it is the reference that governs the office's drafting across all the registers. The manual specifies the typography (the font, the size, the line spacing), the page layout (the margins, the headers, the footers), the heading hierarchy (the heading levels, the numbering, the formatting), the table of contents (the structure, the page numbering), the page numbering (the arabic, the roman, the placement), the front and back matter (the cover, the title page, the copyright, the

acknowledgements), and the plain language standard. The manual is the reference that ensures the office's products are consistent with the government's publication standards, and the manual is the condition of the products' professional presentation.

The government style manual is exemplified by the United States Government Publishing Office Style Manual (the GPO Style Manual, first published in 1894 and revised periodically, with the 2016 edition as the current revision), the Australian Government Style Manual (the Style Manual, first published in 1966 and revised periodically, with the 2020 digital edition as the current revision), and the European Commission's English Style Guide (the Style Guide, first published in 1993 and revised periodically, with the 2022 edition as the current revision). The three manuals are the exemplars of the government style manual, and the geospatial office's drafting follows the manual that the jurisdiction adopts. The manuals share the plain language standard, the professional presentation, and the consistent application, and the manuals differ in the specific conventions (the spelling, the punctuation, the capitalization) that reflect the jurisdiction's linguistic traditions.

The plain language standard is the standard that the government's documents should be written in plain language, and it is the standard that the office's drafting follows across all the registers. The plain language is the language that is clear, concise, and direct, and the plain language is the language that the consumer can understand without the specialized knowledge. The plain language standard is not the simplification of the intelligence; the plain language standard is the clear communication of the intelligence, and the clear communication is the condition of the intelligence's usefulness. The plain language standard is exemplified by the United States Plain Writing Act of 2010 (which requires the federal agencies to use the plain language in the documents intended for the public), the United Kingdom's Government Communication Service's plain language guidance, and the international plain language movement (the Plain Language Association InterNational, the Center for Plain Language).

The plain language standard's application to the geospatial intelligence requires the careful balance between the clarity and the precision. The geospatial intelligence uses the technical terms (the NDVI, the HEC-RAS, the kriging, the weights of evidence) that are precise but not plain, and the officer must decide when to use the technical term (when the audience is the technical specialist who understands the term) and when to use the plain-language equivalent (when the audience is the decision-maker who does not understand the term). The balance is the judgment that the officer exercises in the drafting, and the balance is the condition of the intelligence's accessibility and the precision. The general principle is that the technical term is used when the term is necessary for the precision, with the plain-language explanation provided; the plain-language equivalent is used when the term is not necessary for the precision, with the technical term omitted.

The plain language standard is maintained through the drafting, the editing, and the review. The drafting produces the intelligence in the plain language, with the officer drafting the product with the audience's comprehension in mind. The editing refines the language for the clarity and the concision, with the editor (the office manager or the senior officer) reviewing the draft and suggesting the improvements. The review verifies that the language meets the standard, with the reviewer checking the language's clarity, concision, and directness. The office that maintains the plain language standard produces intelligence that is accessible to its consumers, and the standard is the communication discipline of the geospatial office.

The government style manual is not a constraint on the drafting; the manual is a resource for the drafting. An officer who treats the manual as a constraint produces products that are compliant but uninspired; an officer who treats the manual as a resource produces products that are compliant and effective. The manual provides the framework within which the officer exercises the drafting craft, and the framework is the condition of the craft's consistency. The office that maintains the manual's standards produces products that are professional and consistent, and the standards are the condition of the products' credibility.

35.7 The Cartographic Communication

The cartographic communication is the communication through the maps, and it is the communication that is distinctive to the geospatial profession. The maps are the geospatial office's primary communication medium, and the maps' design determines the intelligence's communicability. The cartographic communication comprises the map design (the symbology, the generalization, the layout), the map production (the cartographic software, the format, the quality), and the map evaluation (the readability, the accuracy, the effectiveness).

The map design is the cartographic communication's first component, and it is the component that determines the map's communicability. The symbology is the set of the symbols, the colours, and the patterns that represent the geographic features, and the symbology must be clear (the symbols are distinguishable), consistent (the same feature is represented by the same symbol throughout the map), and meaningful (the symbol's appearance conveys the feature's characteristics). The generalization is the simplification of the geographic features for the map's scale, and the generalization must be appropriate (the features are simplified to the level that the scale allows), consistent (the same feature type is generalized by the same rules), and honest (the generalization does not misrepresent the features). The layout is the arrangement of the map's components (the map area, the legend, the title, the scale bar, the north arrow, the classification marking), and the layout must be balanced (the components are arranged for the visual balance), clear (the components are distinguishable), and professional (the layout follows the government style manual's standards).

The map production is the cartographic communication's second component, and it is the component that produces the map in the formats the consumers require. The cartographic software is the software that produces the map, with the office's product studio (treated in Chapter 29, section 29.6) using the ESRI ArcGIS, the QGIS, or the custom-developed software. The format is the file format that the map is delivered in, with the GeoTIFF for the raster products, the Shapefile for the vector products, the PDF for the printed products, and the KML for the virtual globe viewers. The quality is the map's visual quality, with the resolution (the dots per inch for the printed products, the pixel size for the digital products), the colour (the colour model: the RGB for the digital, the CMYK for the printed), and the file size (the balance between the quality and the download time).

The map evaluation is the cartographic communication's third component, and it is the component that verifies the map's effectiveness. The readability is the map's ability to be read and understood by the consumer, with the readability assessed through the consumer testing (the consumer's ability to find the information, interpret the symbols, and understand the message). The accuracy is the map's ability to represent the geographic features correctly, with the accuracy assessed through the validation (the comparison with the ground truth). The effectiveness is the map's ability

to communicate the intelligence, with the effectiveness assessed through the consumer's feedback (the consumer's use of the map for the decisions).

The cartographic communication's discipline is the discipline of producing the maps that are clear, accurate, and effective, and the discipline is developed through the experience (the production of the multiple maps), the study (the examination of the comparative maps), and the mentorship (the guidance of the senior officers). The discipline is the cartographic dimension of the geospatial profession, and the discipline is the dimension that the training (treated in Chapter 31, section 31.3, the product studio course) and the judgment (treated in Chapter 34) support. The cartographic communication is the condition of the maps' communicability, and the communicability is the condition of the intelligence's use.

35.8 The Briefing and the Oral Communication

The briefing and the oral communication are the communication through the spoken word, and they are the communication that complements the written products. The geospatial officer communicates orally in the briefings (the event briefings, the seasonal briefings, the annual briefings), the presentations (the inter-agency presentations, the conference presentations, the public presentations), and the testimony (the regulatory hearings, the audit responses, the parliamentary testimony). The oral communication is the communication that the written products cannot fully replace, because the oral communication allows the real-time interaction, the clarification of the ambiguities, and the persuasion of the decision-makers.

The briefing is the oral communication that the officer provides to the consumers, and it is the communication that supports the consumers' decisions. The event briefing is the briefing provided during the event, with the office manager or the liaison officer briefing the emergency management authority's emergency operations centre on the office's intelligence (treated in Chapter 24, section 24.9). The seasonal briefing is the briefing provided at the seasonal milestones, with the resource economics specialist briefing the strategic reserve authority's procurement director on the yield forecast and the procurement advisory (treated in Chapter 25, section 25.9). The annual briefing is the briefing provided at the annual review, with the office manager briefing the desk chief and the bureau chief on the office's annual performance. The briefing's structure comprises the briefing's purpose (the briefing's objective), the briefing's content (the intelligence presented), the briefing's visual aids (the maps, the charts, the slides), and the briefing's interaction (the questions and the answers).

The presentation is the oral communication that the officer provides to the broader audience, and it is the communication that shares the office's work with the profession and the public. The inter-agency presentation is the presentation to the other government agencies, with the office manager or the technical lead presenting the office's work at the inter-agency forums. The conference presentation is the presentation to the professional community, with the specialists presenting the office's research at the professional conferences (treated in Chapter 31, section 31.6). The public presentation is the presentation to the public, with the office manager or the specialists presenting the office's work at the public forums (the open days, the university lectures, the media interviews). The presentation's structure comprises the presentation's purpose, the presentation's content, the presentation's visual aids, and the presentation's delivery (the voice, the pace, the body language).

The testimony is the oral communication that the officer provides to the formal proceedings, and it is the communication that supports the legal and the regulatory processes. The regulatory hearing is the testimony before the sectoral regulator's hearing (treated in Chapter 26, section 26.9), with the hazard specialist testifying on the office's monitoring intelligence and the compliance verification report. The audit response is the testimony before the national audit authority, with the office manager testifying on the office's compliance with the financial, the operational, and the security standards. The parliamentary testimony is the testimony before the parliamentary committee, with the office manager or the senior officers testifying on the office's work and the government's geospatial intelligence. The testimony's structure comprises the testimony's purpose, the testimony's content (the written statement submitted in advance), the testimony's delivery (the sworn testimony, the questioning by the committee), and the testimony's documentation (the transcript, the record).

The oral communication's discipline is the discipline of producing the briefings, the presentations, and the testimony that are clear, accurate, and effective, and the discipline is developed through the experience, the study, and the mentorship. The discipline comprises the preparation (the briefing's preparation, the presentation's rehearsal, the testimony's preparation), the delivery (the clear speech, the appropriate pace, the effective body language), and the interaction (the listening to the questions, the clear answers, the honest acknowledgment of the limitations). The oral communication is the condition of the intelligence's communication to the consumers who do not read the written products, and the communication is the condition of the intelligence's use by the broader audience.

35.9 Worked Example: Drafting the Procurement Advisory for the Northfields Drought

The drafting is illustrated through the procurement advisory for the Northfields drought, treated in Chapter 25, section 25.9, and Chapter 20, section 20.9. The advisory is the office's forward-sensor product, drafted in the analytical register, and delivered to the strategic reserve authority on 22 July 2024. The worked example traces the advisory's drafting through the four sections (the yield forecast, the stock level assessment, the supply-demand analysis, and the release or buy signal), demonstrating the analytical register's execution and the drafting discipline's application.

The advisory's drafting begins on 22 July 2024 at 12:00, with the resource economics specialist Mr. Jonas Adekoya (treated in Chapter 31, section 31.9) compiling the advisory's draft from the completed analysis. Mr. Adekoya opens the advisory's template (the standard 4-to-6-page PDF template, with the Internal classification marking, the header, and the footer) and begins with the yield forecast section.

The yield forecast section is drafted as follows: "The Northfields' wheat yield is forecast at 2.1 tonnes per hectare, with the 90-per-cent confidence interval of 1.8 to 2.4 tonnes per hectare. The forecast is based on the WOFOST crop growth model, configured with the observed conditions (the NDVI at the 22nd percentile of the historical distribution, the NDWI at the 18th percentile, the soil moisture from the NASA SMAP satellite at 25 per cent below normal), the weather to date (the March-to-July rainfall at 65 per cent of normal, the temperatures at 1.2°C above normal), and the weather outlook for the remaining growing period (the meteorological service's 30-day forecast predicts the below-normal rainfall and the above-normal temperatures to continue). The forecast is a

19-per-cent downward revision from the planting forecast of 2.6 tonnes per hectare, reflecting the drought's intensification between the planting and the heading. The forecast's confidence interval reflects the drought's intensity and the weather outlook's uncertainty. The forecast was validated by the statistical regression model (the NDVI at the heading stage and the historical yield), which produced a forecast of 2.0 tonnes per hectare, within the WOFOST model's confidence interval." The section is thorough (the forecast, the assumptions, the uncertainties, the validation are all presented), balanced (the uncertainties are honestly presented), and honest about the uncertainties (the confidence interval's basis is explained).

The stock level assessment section is drafted as follows: "The reserve's wheat stocks are 1.85 million tonnes as of 22 July 2024, distributed across the four storage depots (the Northfields Depot 0.65 million tonnes, the Portsmith Depot 0.50 million tonnes, the Capital Depot 0.40 million tonnes, the Western Depot 0.30 million tonnes). The stock-to-consumption ratio is 3.7 months (the stocks represent 3.7 months of the national wheat consumption of 0.50 million tonnes per month), below the authority's target of 4.5 months. The stock levels have declined by 12 per cent over the past 12 months, reflecting the higher-than-expected consumption (the 2023-2024 consumption was 2 per cent above the forecast) and the lower-than-expected procurement (the 2023 harvest was 5 per cent below the 2023 forecast). The stock levels' source is the reserve authority's data-sharing arrangement with the office, with the data provided weekly through the secure data interface." The section is thorough (the stocks, the ratio, the trend, the source are all presented) and balanced (the trend's causes are honestly presented).

The supply-demand analysis section is drafted as follows: "The production estimate is 2.52 million tonnes (the yield of 2.1 tonnes per hectare multiplied by the planted area of 1.20 million hectares). The consumption estimate is 0.60 million tonnes per year (the population of 12 million multiplied by the per-capita consumption rate of 50 kilograms per person per year). The carry-over from the previous season's stocks is 0.35 million tonnes. The annual supply is 2.87 million tonnes (the production of 2.52 plus the carry-over of 0.35). The annual supply-demand balance is a surplus of 2.27 million tonnes (the supply of 2.87 minus the consumption of 0.60). The reserve's projected level after the harvest is 2.10 million tonnes (the current stocks of 1.85 minus the consumption until the harvest of 0.15 plus the procurement from the harvest of 0.40), representing 4.2 months of consumption, below the 4.5-month target." The section is thorough (the production, the consumption, the balance, the projected reserve are all presented) and balanced (the balance's implications are honestly presented).

The release or buy signal section is drafted as follows: "The advisory recommends a buy of 0.20 million tonnes, to be executed before the harvest (by 31 August 2024). The buy is recommended because the drought-reduced yield (2.1 tonnes per hectare, 19 per cent below the planting forecast) and the below-target stock-to-consumption ratio (3.7 months, below the 4.5-month target) warrant the precautionary procurement to maintain the reserve's adequacy through the next consumption cycle. The buy should be executed before the harvest to avoid the price increase that the drought-reduced supply will cause after the harvest (the historical pattern suggests an 8-to-12-per-cent price increase following the drought-reduced harvests). The buy's fiscal cost is estimated at $80 million (the 0.20 million tonnes at the prevailing international wheat price of $400 per tonne), within the

authority's procurement budget." The section is direct (the signal and the quantity are clearly stated) and thorough (the rationale, the timing, and the fiscal cost are all presented).

The advisory's review is conducted by the office manager on 22 July 2024 at 14:30, with the review verifying the advisory's accuracy (the yield forecast, the stock levels, the supply-demand analysis, and the signal are correct), the advisory's thoroughness (the four sections are present and the analysis is complete), the advisory's balance (the uncertainties are honestly presented, the confidence interval is explained, the trend's causes are identified), and the advisory's classification (the Internal marking is applied to the header, the footer, and the metadata). The office manager's review identifies two minor issues: (1) the yield forecast section's first sentence should explicitly state the forecast's date ("as of 22 July 2024"), and (2) the stock level assessment section should note the data's vintage ("the data as of 21 July 2024, provided by the reserve authority on 22 July 2024"). The office manager communicates the issues to Mr. Adekoya, who makes the corrections, and the office manager authorizes the release at 15:00.

The advisory's delivery is conducted by the production operator at 15:15 on 22 July 2024, with the advisory delivered to the reserve authority's procurement director through the internal distribution channel. The delivery is acknowledged by the procurement director at 16:30, and the advisory's release is logged in the office's distribution log. The advisory's archive copy is retained in the office's classified archive, with the 10-year retention period (the standard retention for the procurement advisories, treated in Chapter 12, section 12.4).

The advisory's outcome is documented in the decision feedback received from the reserve authority on 25 July 2024 (treated in Chapter 25, section 25.9): the authority decided to buy 0.20 million tonnes, as recommended by the advisory, and the buy was executed before the harvest. The advisory's forecast was accurate (the harvest report published in September 2024 reported the Northfields' wheat yield at 2.05 tonnes per hectare, within the advisory's confidence interval of 1.8 to 2.4 tonnes per hectare), and the advisory's signal was correct (the buy maintained the reserve's adequacy and avoided the 8-per-cent price increase that occurred after the harvest). The advisory's outcome is the illustration of the analytical register's effectiveness, and the outcome is the condition of the office's credibility with the reserve authority.

This worked example illustrates the drafting's full operation: the yield forecast section's drafting (the thorough presentation of the forecast, the assumptions, the uncertainties, the validation), the stock level assessment section's drafting (the thorough presentation of the stocks, the ratio, the trend, the source), the supply-demand analysis section's drafting (the thorough presentation of the production, the consumption, the balance, the projected reserve), the release or buy signal section's drafting (the direct presentation of the signal and the thorough presentation of the rationale, the timing, and the fiscal cost), the office manager's review (the accuracy, the thoroughness, the balance, the classification, the two minor corrections), the delivery (the internal distribution channel, the acknowledgement, the archive), and the outcome (the authority's decision matching the advisory's signal, the forecast's accuracy, the signal's correctness). The analytical register's discipline was maintained throughout: the advisory was thorough (the four sections, the complete analysis), balanced (the uncertainties honestly presented), and honest about the uncertainties (the confidence interval explained, the trend's causes identified). The drafting discipline is the condition of the advisory's usefulness for the reserve authority's decision, and the discipline is the condition of the

office's credibility. The procurement advisory for the Northfields drought is the illustration of the analytical register's execution, and the advisory is the example that this chapter provides for the drafting and communication craft. The craft is the communication dimension of the geospatial profession, and the dimension is the condition of the office's intelligence delivery, and the delivery is the condition of the office's contribution to the government's decisions.

Chapter 36 — The Geospatial Officer's Toolkit

36.1 The Toolkit's Purpose

The geospatial officer's toolkit comprises the reference works, the software, the data sources, the professional networks, the standards bodies, the continuing education, and the professional development path that support the work over the career. The toolkit is not a list of tools but a description of the professional resources on which the officer draws, stated in the order in which the officer encounters them over a career. The toolkit's purpose is to provide the officer with the resources that the doctrine, the procedures, and the job breakdown sheets treated throughout this Handbook require, and the toolkit is the condition of the officer's ability to execute the doctrine, the procedures, and the job breakdown sheets.

The toolkit is distinguished from the information system (treated in Chapter 29) by its scope and its character. The information system is the technical infrastructure that the office operates, comprising the six submodules and the integration with the core system. The toolkit is the professional infrastructure that the officer draws on, comprising the resources that the officer brings to the information system's operation. The information system is the office's collective resource, maintained by the office and shared by all officers. The toolkit is the officer's individual resource, built by the officer over the career and carried by the officer across the roles and the offices. The information system and the toolkit are complementary: the information system provides the capabilities, and the toolkit provides the competence to use the capabilities.

The toolkit's seven components are the reference works, the software, the data sources, the professional networks, the standards bodies, the continuing education, and the professional development path. Each component is examined in this chapter, with the examination covering the component's content, the component's acquisition, and the component's maintenance. The seven components together constitute the toolkit, and the toolkit is the professional foundation of the geospatial officer's career.

The toolkit is built over the career, with the officer acquiring the components progressively. The reference works are acquired first, through the university education and the foundational training. The software is acquired next, through the submodule courses and the on-the-job training. The data sources are acquired through the office's data-sharing arrangements and the officer's familiarity with the data landscape. The professional networks are acquired through the office's inter-agency work and the officer's professional engagement. The standards bodies are acquired through the officer's participation in the standards' development and the standards' application. The continuing education is acquired throughout the career, through the refresher courses, the update courses, and the advanced courses. The professional development path is acquired through the career's progression, with the path providing the framework for the toolkit's development. The seven components' progressive acquisition is the reason the toolkit is the career's product rather than the appointment's gift, and the progressive acquisition is the reason the toolkit is the experienced officer's resource rather than the novice's.

36.2 The Reference Works

The reference works are the first component of the toolkit, because the reference works are the resources that the officer consults from the first day of the career to the last. The reference works comprise the scientific literature, the technical documentation, the standards, and the professional publications. Each category serves a different purpose in the officer's work, and the four categories together provide the officer with the knowledge foundation that the work requires.

The scientific literature is the peer-reviewed literature of the geospatial sciences, including the hydrology, the remote sensing, the photogrammetry, the cartography, and the geographic information science. The literature comprises the journals (the Remote Sensing of Environment, the ISPRS Journal of Photogrammetry and Remote Sensing, the International Journal of Applied Earth Observation and Geoinformation, the Journal of Hydrology, the Hydrology and Earth System Sciences, the Cartography and Geographic Information Science, the International Journal of Geographical Information Science), the books (the textbooks and the monographs that the officer consults for the foundational knowledge and the specialized topics), and the conference proceedings (the ISPRS Congress, the IGARSS, the AGU Fall Meeting, the EGU General Assembly). The officer consults the literature for the methods, the case studies, and the comparative practice, and the literature is the mechanism by which the officer's work is informed by the broader scientific community.

The technical documentation is the documentation of the software and the systems the office uses, including the processing engine (the documentation of the ENVI, the ERDAS Imagine, the Sentinel Toolbox, the GDAL), the modeling workbench (the documentation of the HEC-HMS, the HEC-RAS, the WOFOST, the DSSAT, the MaxEnt, the ArcSDM), the product studio (the documentation of the ArcGIS, the QGIS, the GeoServer), and the catalog and archive (the documentation of the PostGIS, the GeoNetwork, the CKAN). The officer consults the documentation for the software's operation, the parameters' configuration, and the troubleshooting, and the documentation is the mechanism by which the officer operates the software correctly and efficiently.

The standards are the international and national standards that govern the geospatial work, including the ISO geospatial standards (the ISO 19100 series, comprising the ISO 19115 for the metadata, the ISO 19119 for the services, the ISO 19136 for the Geography Markup Language, the ISO 19142 for the Web Feature Service), the Open Geospatial Consortium standards (the OGC Web Map Service, the Web Feature Service, the Web Coverage Service, the Catalog Service for the Web, the Sensor Observation Service), and the national mapping standards (the Federal Geographic Data Committee standards in the United States, the Intergovernmental Committee on Surveying and Mapping standards in Australia, the Ordnance Survey standards in the United Kingdom). The officer consults the standards for the data's structure, the services' interfaces, and the quality's requirements, and the standards are the mechanism by which the officer's work is interoperable with the broader geospatial community.

The professional publications are the publications of the professional associations, including the journals (the Photogrammetric Engineering and Remote Sensing by the ASPRS, the Journal of Spatial Science by the SSSI, the Survey Review by the Commonwealth Association of Surveying and Land Economy), the magazines (the GeoConnexion, the GIS Lounge, the Directions Magazine), and the newsletters (the ISPRS Newsletter, the OGC Newsletter, the GEO Newsletter). The officer

consults the publications for the profession's news, the trends, and the opportunities, and the publications are the mechanism by which the officer stays current with the profession's development.

36.3 The Software

The software is the second component of the toolkit, and it is the software that the officer operates on a daily basis. The software comprises the geospatial information system, the remote sensing processing software, the statistical analysis software, the cartographic production software, and the office productivity software. Each category supports a different aspect of the officer's work, and the five categories together provide the officer with the software capability that the work requires.

The geospatial information system is the system that hosts the office's pipeline, treated in Chapter 29, and it is the system that the officer operates through the six submodules. The GIS comprises the commercial software (the ESRI ArcGIS Pro, the ArcGIS Enterprise, the ArcGIS Online) and the open-source software (the QGIS, the PostGIS, the GeoServer, the GeoNode). The officer's proficiency in the GIS includes the data management (the ingestion, the storage, the indexing), the spatial analysis (the overlay, the buffer, the network analysis), the cartographic production (the map design, the symbology, the layout), and the publication (the web services, the web maps, the web apps). The GIS proficiency is developed through the submodule courses (treated in Chapter 31, section 31.3) and the on-the-job training.

The remote sensing processing software is the software that the processing engine uses for the georeferencing, the calibration, the mosaicking, the projection management, the cloud masking, and the data fusion (treated in Chapter 9). The software comprises the commercial software (the ENVI by L3Harris Geospatial, the ERDAS Imagine by Hexagon Geospatial, the eCognition by Trimble) and the open-source software (the Sentinel Toolbox by ESA, the GRASS GIS, the GDAL, the Orfeo ToolBox). The officer's proficiency in the remote sensing software includes the imagery display (the band combinations, the contrast enhancement, the false colour), the image processing (the georeferencing, the calibration, the mosaicking), the image classification (the supervised, the unsupervised, the object-based), and the change detection (the image differencing, the post-classification comparison).

The statistical analysis software is the software that the modeling workbench uses for the statistical analysis, the model fitting, and the validation. The software comprises the general-purpose statistical software (the R, the Python with the NumPy, the SciPy, the pandas, the scikit-learn), the specialized geospatial statistical software (the GeoDa, the spdep, the gstat), and the domain-specific software (the HEC-HMS and the HEC-RAS for the hydrology, the DSSAT and the WOFOST for the crop modelling, the MaxEnt and the ArcSDM for the prospectivity). The officer's proficiency in the statistical software includes the data manipulation (the data cleaning, the data transformation, the data aggregation), the statistical analysis (the descriptive statistics, the hypothesis testing, the regression), and the model fitting (the calibration, the validation, the uncertainty analysis).

The cartographic production software is the software that the product studio uses for the cartographic design, the symbology, and the format conversion (treated in Chapter 29, section 29.6). The software comprises the GIS-integrated cartographic software (the ArcGIS Pro, the QGIS), the specialized cartographic software (the Adobe Illustrator, the Adobe InDesign, the CorelDraw for the high-quality map production), and the web cartographic software (the Mapbox, the Carto, the

Leaflet). The officer's proficiency in the cartographic software includes the map design (the layout, the symbology, the generalization), the map production (the format conversion, the quality control, the output), and the web cartography (the web map design, the interactive maps, the story maps).

The office productivity software is the software that the officer uses for the communication, the reporting, and the administration, including the word processing (the Microsoft Word, the LibreOffice Writer), the spreadsheet (the Microsoft Excel, the LibreOffice Calc), the presentation (the Microsoft PowerPoint, the LibreOffice Impress), and the email (the Microsoft Outlook, the Mozilla Thunderbird). The officer's proficiency in the office productivity software includes the document production (the reports, the briefings, the procedures), the data analysis (the spreadsheets, the pivot tables, the charts), the presentation (the slides, the visual aids, the delivery), and the communication (the email, the calendar, the tasks). The office productivity software is the software that every officer uses daily, and the proficiency is the condition of the officer's administrative efficiency.

36.4 The Data Sources

The data sources are the third component of the toolkit, and they are the sources from which the office acquires the data for its pipeline. The data sources comprise the satellite imagery, the aerial imagery, the sensor feeds, the base layers, and the thematic data. Each category provides a different type of the data, and the five categories together provide the office with the data that the pipeline requires.

The satellite imagery is the imagery from the Earth observation satellites, including the free-and-open imagery (the Landsat 8 and 9 by the USGS and NASA, the Sentinel-1, 2, 3, 5P by the ESA Copernicus programme, the MODIS by the NASA Terra and Aqua) and the commercial imagery (the Maxar WorldView-1, 2, 3, 4, the GeoEye, the QuickBird, the Airbus Pleiades, Pleiades Neo, the SPOT, the Planet PlanetScope, the SkySat), treated in Chapter 7. The officer's familiarity with the satellite imagery includes the sensors' characteristics (the spatial resolution, the spectral resolution, the temporal resolution, the radiometric resolution), the tasking and the acquisition (the tasking APIs, the archive search, the pricing), and the data's processing level (the raw, the georeferenced, the orthorectified, the surface reflectance).

The aerial imagery is the imagery from the aerial photography and the airborne remote sensing, including the aerial photography (the vertical photography, the oblique photography, the digital aerial cameras), the LiDAR (the airborne laser scanning, with the point density, the vertical accuracy, the classification), and the airborne geophysics (the aeromagnetic, the radiometric, the electromagnetic). The aerial imagery is the high-resolution data that complements the satellite imagery, with the aerial photography providing the high-resolution optical coverage, the LiDAR providing the high-resolution elevation, and the airborne geophysics providing the subsurface characterization. The officer's familiarity with the aerial imagery includes the acquisition (the flight planning, the sensor configuration), the processing (the photogrammetry, the point cloud processing, the geophysical processing), and the application (the base layer production, the change detection, the prospectivity).

The sensor feeds are the continuous data streams from the ground-based sensors, including the rain gauges (the tipping-bucket, the weighing, the optical), the river stage sensors (the pressure, the radar, the ultrasonic), the weather stations (the temperature, the humidity, the wind, the solar

radiation), the soil moisture sensors (the time-domain reflectometry, the capacitance), and the air quality monitors (the particulate matter, the gases). The sensor feeds are the real-time data that the office uses for the event-driven intelligence (the flood modelling, the crop monitoring), and the officer's familiarity with the sensor feeds includes the sensor types, the data formats (the OneRain, the ALERT2, the HYDRA), and the data access (the APIs, the web services, the data-sharing arrangements).

The base layers are the maintained datasets that provide the geographic context, including the digital elevation model (the SRTM, the ASTER GDEM, the national LiDAR DEM), the land cover (the ESA WorldCover, the USGS National Land Cover Database, the national land cover), the hydrography (the HydroSHEDS, the national hydrography datasets), the administrative boundaries (the FAO GAUL, the national administrative boundaries), and the cadastre feed (from the cadastral authority, treated in Chapter 23), treated in Chapter 8. The base layers are the foundation of the office's spatial data infrastructure, and the officer's familiarity with the base layers includes the data's coverage, the data's accuracy, the data's currency, and the data's access.

The thematic data is the domain-specific data, including the geological survey data (the geological maps, the geochemical data, the geophysical data, treated in Chapter 22), the agricultural statistics (the cropped area, the yield, the production, treated in Chapter 20), the infrastructure asset registries (the energy infrastructure, the transportation infrastructure, treated in Chapter 21), and the demographic data (the census, the population estimates, treated in Chapter 28). The thematic data is the data that the office's domain plug-ins consume, and the officer's familiarity with the thematic data includes the data's sources, the data's formats, the data's quality, and the data's access.

36.5 The Professional Networks

The professional networks are the fourth component of the toolkit, and they are the networks of the professional colleagues and the professional associations that support the officer's work. The professional networks comprise the office's internal network, the government's inter-agency network, the national professional associations, and the international professional associations. Each network provides a different type of the support, and the four networks together provide the officer with the professional community that the work requires.

The office's internal network is the network of the office's officers, who share the knowledge and the experience through the daily work. The internal network comprises the office manager (the mentor and the leader), the technical lead (the technical reference), the senior specialists (the domain references), the junior officers (the peers), and the operators (the operational colleagues). The internal network is the officer's first and most immediate professional network, and the network is the mechanism by which the officer learns the office's specific practices, the office's specific tools, and the office's specific consumers. The internal network's strength is the condition of the office's collective competence, and the network is maintained through the daily work, the formal meetings, and the informal interactions.

The government's inter-agency network is the network of the geospatial offices across the government, who share the data and the practice through the inter-agency frameworks. The inter-agency network comprises the emergency management authority (treated in Chapter 24), the strategic reserve authority (treated in Chapter 25), the sectoral regulators (treated in Chapter 26), the water

authority (treated in Chapter 27), the national statistics agency (treated in Chapter 28), and the cadastral authority (treated in Chapter 23). The inter-agency network is the officer's second professional network, and the network is the mechanism by which the officer coordinates the inter-agency work, the data sharing, and the joint production. The inter-agency network's strength is the condition of the government's geospatial coordination, and the network is maintained through the inter-agency frameworks, the coordination meetings, and the joint exercises.

The national professional associations are the associations of the geospatial professionals in the jurisdiction, including the national geospatial societies (the American Society for Photogrammetry and Remote Sensing [ASPRS] in the United States, the Surveying and Spatial Sciences Institute [SSSI] in Australia, the Royal Institution of Chartered Surveyors [RICS] in the United Kingdom, the Deutsche Gesellschaft für Geodäsie, Geoinformation und Landmanagement [DVW] in Germany, the Association Française de Topographie [AFT] in France), the national remote sensing societies (the IEEE Geoscience and Remote Sensing Society, the Canadian Remote Sensing Society), and the national cartographic societies (the British Cartographic Society, the Swiss Society of Cartography). The national associations are the officer's third professional network, and the network is the mechanism by which the officer engages with the national profession, the certifications, and the conferences.

The international professional associations are the associations of the geospatial professionals internationally, including the International Society for Photogrammetry and Remote Sensing (ISPRS), the International Cartographic Association (ICA), the International Federation of Surveyors (FIG), the Open Geospatial Consortium (OGC), the Group on Earth Observations (GEO), and the Committee on Earth Observation Satellites (CEOS). The international associations are the officer's fourth professional network, and the network is the mechanism by which the officer engages with the international profession, the standards, and the cooperation. The international network is particularly important for the officers who participate in the international cooperation (the Copernicus programme, the GEO, the International Charter on Space and Major Disasters), and the network is the condition of the officer's international engagement.

The professional networks are maintained through the active participation, with the officer attending the meetings, contributing to the associations' work, and engaging with the colleagues. The networks are not the passive memberships but the active relationships, and the active relationships are the condition of the networks' value. An officer who maintains the professional networks has access to the knowledge and the support of the profession; an officer who does not is isolated, with the isolation limiting the officer's access to the comparative practice, the emerging methods, and the professional opportunities.

36.6 The Standards Bodies

The standards bodies are the fifth component of the toolkit, and they are the bodies that develop and maintain the standards that govern the geospatial work. The standards bodies comprise the international standards bodies, the national standards bodies, and the professional standards bodies. Each type of the standards body develops a different category of the standards, and the three types together provide the standards framework that the geospatial work requires.

The international standards bodies include the International Organization for Standardization (ISO), the Open Geospatial Consortium (OGC), and the International Federation of Surveyors (FIG). The ISO develops the ISO 19100 series of the geographic information standards, comprising the ISO 19115 (the metadata), the ISO 19119 (the services), the ISO 19136 (the Geography Markup Language), the ISO 19142 (the Web Feature Service), and the other standards that govern the geographic information's structure, services, and quality. The OGC develops the open standards for the geospatial web services, comprising the Web Map Service (WMS), the Web Feature Service (WFS), the Web Coverage Service (WCS), the Catalog Service for the Web (CSW), and the Sensor Observation Service (SOS). The FIG develops the surveying standards, comprising the standards for the cadastre, the land administration, and the hydrography.

The national standards bodies include the national mapping agencies (the USGS in the United States, the Ordnance Survey in the United Kingdom, the IGN in France, the IGN Spain, the Kadaster in the Netherlands), the national geological surveys (the USGS, the BGS in the United Kingdom, the BRGM in France, the BGR in Germany), and the national standards organizations (the ANSI in the United States, the BSI in the United Kingdom, the AFNOR in France, the DIN in Germany, the Standards Australia). The national standards bodies develop the national standards that supplement the international standards, with the national standards addressing the jurisdiction's specific requirements (the national coordinate systems, the national data formats, the national quality standards).

The professional standards bodies include the professional associations that develop the professional standards, including the certification standards (the ASPRS certifications, the SSSI certifications, the RICS certifications, the GIS Certification Institute's GISP certification) and the ethical standards (the ASPRS Code of Ethics, the SSSI Code of Ethics, the RICS Rules of Conduct). The professional standards are the standards that govern the officer's professional conduct, and the standards are the condition of the officer's professional standing.

The standards bodies are the resources that the officer consults for the standards that govern the work, and the standards bodies are the resources that the officer participates in for the development of the standards. The officer's consultation of the standards bodies includes the standards' review (the officer reads the standards before applying them), the standards' application (the officer applies the standards in the daily work), and the standards' interpretation (the officer interprets the standards for the specific situations). The officer's participation in the standards bodies includes the standards' development (the officer contributes to the standards' drafting and the revision through the national mirror committees or the OGC working groups), the standards' promotion (the officer promotes the standards' adoption in the office and the inter-agency), and the standards' governance (the officer participates in the standards bodies' governance through the voting or the board membership).

The standards bodies are the institutional framework for the geospatial profession's standards, and the framework is the condition of the profession's consistency and quality. An officer who participates in the standards bodies contributes to the development of the standards and stays current with the standards' evolution; an officer who does not participate is a passive consumer of the standards, with the passivity limiting the officer's influence on the standards and the officer's currency with the standards' evolution. The participation is the mechanism by which the officer

contributes to the profession's standards, and the participation is the condition of the profession's collective advancement.

36.7 The Continuing Education and the Professional Development Path

The continuing education and the professional development path are the sixth and the seventh components of the toolkit, and they are the components that maintain and advance the officer's competence over the career. The continuing education is the ongoing education that the officer completes throughout the career, treated in Chapter 31, section 31.6. The professional development path is the path from the junior operator to the office manager, treated in Chapter 4 and Chapter 30, section 30.4. The two components are complementary: the continuing education maintains the officer's current competence, and the professional development path advances the officer's future competence.

The continuing education comprises the refresher courses, the update courses, the advanced courses, and the self-directed learning. The refresher courses refresh the officer's competence in the foundational, the submodule, and the domain courses, and the courses are delivered every 2 to 3 years. The update courses update the officer's competence in the new methods, the new tools, and the new procedures, and the courses are delivered as the new methods, the new tools, and the new procedures are introduced. The advanced courses advance the officer's competence in the specialized topics, the emerging methods, and the leadership skills, and the courses are delivered to the senior officers. The self-directed learning is the learning that the officer conducts through the online courses, the professional literature, and the conference attendance, and the learning is the officer's individual contribution to the continuing education.

The continuing education is not optional; it is a requirement of the officer's certification, and the certification must be maintained through the continuing education (the 120 continuing education hours over the 3-year recertification cycle, treated in Chapter 31, section 31.7). An officer who does not complete the continuing education loses the certification, and the loss of the certification means the officer is no longer qualified to operate the submodule or the domain. The continuing education is the institutional mechanism for the officer's ongoing professional development, and the mechanism is the condition of the officer's competence over the career.

The professional development path is the path from the junior operator to the office manager, comprising the three principal career paths (the pipeline path, the production path, and the domain path, treated in Chapter 30, section 30.4), the lateral moves between the paths, and the convergence at the technical lead and the office manager ranks. The path is the institutional framework for the officer's career, and the path is the condition of the officer's advancement. The path is not a predetermined trajectory; the path is a set of opportunities and the officer's choices among them. An officer who follows the path develops the expertise that the path provides, and the expertise is the foundation for the officer's advancement. An officer who does not follow a path develops the expertise haphazardly, with the haphazardness limiting the officer's advancement.

The professional development path and the toolkit are mutually reinforcing. The path provides the framework for the toolkit's development, with each stage of the path requiring the acquisition of the specific toolkit components (the foundational stage requires the reference works and the software basics, the operational stage requires the software proficiency and the data sources familiarity, the

advanced stage requires the professional networks and the standards bodies engagement, the leadership stage requires the continuing education's leadership components and the professional development path's leadership aspects). The toolkit provides the resources for the path's progression, with the toolkit's components enabling the officer to perform at each stage of the path and to advance to the next stage. The mutual reinforcement is the reason the path and the toolkit are the complementary components of the officer's career, and the mutual reinforcement is the condition of the officer's professional growth.

The continuing education and the professional development path are the conditions of the officer's lifelong learning. The geospatial profession is the evolving profession, with the new methods, the new tools, and the new applications emerging continuously, and the officer's competence must evolve with the profession. The continuing education maintains the officer's currency, and the professional development path advances the officer's capability. The two together constitute the officer's lifelong learning, and the lifelong learning is the condition of the officer's competence over the career, and the competence over the career is the condition of the officer's contribution to the office's work and the profession's advancement.

36.8 The Toolkit and the Handbook

The toolkit and the Handbook are related, because the Handbook is the resource that the officer consults for the doctrine, the procedures, and the job breakdown sheets, and the toolkit is the resource that the officer draws on for the reference works, the software, the data sources, the professional networks, the standards bodies, the continuing education, and the professional development path. The Handbook provides the rules and the judgment (treated in Chapter 34, section 34.8), and the toolkit provides the resources that the rules and the judgment require. The Handbook and the toolkit are complementary, and the complementarity is the condition of the officer's competence.

The Handbook's doctrine, treated throughout the 35 preceding chapters, provides the officer with the profession's collective knowledge. The doctrine comprises the foundational principles (the commercial firewall, the pipeline architecture, the shared-backbone and the swappable-heads design, treated in Part I), the pipeline capabilities (the base layer, the acquisition, the processing, the modeling, the production, the archive, treated in Part II), the role family and the job breakdown sheets (the six seats and their daily work, treated in Part III), the four domains (the flood risk, the food security, the energy infrastructure, the minerals, treated in Part IV), the sectoral interfaces (the emergency management authority, the strategic reserve authority, the sectoral regulator, the water authority, the cadastral authority, the national statistics agency, treated in Part V), the methods and tools (the information system, the role family and qualifications, the training, the continuous improvement, treated in Part VI), and the comparative practice and craft (the comparative traditions, the judgment, the drafting and communication, treated in Part VII). The doctrine is the profession's collective knowledge, and the doctrine is the resource that the officer consults for the rules.

The toolkit's seven components, treated in this chapter, provide the officer with the profession's collective resources. The reference works provide the knowledge, the software provides the capability, the data sources provide the input, the professional networks provide the community, the standards bodies provide the framework, the continuing education provides the currency, and the

professional development path provides the advancement. The seven components are the profession's collective resources, and the resources are the conditions that the officer draws on for the work.

The relationship between the Handbook and the toolkit is the relationship between the knowledge and the resources. The knowledge without the resources is the theory that cannot be applied, and the resources without the knowledge are the tools that cannot be used effectively. The knowledge and the resources together constitute the officer's competence, and the competence is the condition of the officer's work. The Handbook is the knowledge, and the toolkit is the resources, and the two together are the geospatial officer's professional foundation.

The Handbook's completion with this chapter is the completion of the doctrine, and the completion is the culmination of the work that began with the foundational principles in Chapter 1. The 36 chapters together constitute the doctrine, and the doctrine is the profession's collective knowledge, codified for the officer's use. The toolkit, treated in this final chapter, is the resource that the officer draws on for the work, and the toolkit is the condition of the doctrine's application. The Handbook and the toolkit together are the geospatial officer's professional foundation, and the foundation is the condition of the officer's competence, the office's effectiveness, and the profession's contribution to the government's geospatial intelligence.

36.9 Worked Example: The Toolkit of Dr. Maren Holloway

The toolkit's deployment is illustrated through the example of Dr. Maren Holloway, the hazard specialist whose recruitment and career progression are treated in Chapter 30, section 30.9, and whose judgment during the Riverbourne Basin flood event is treated in Chapter 34, section 34.9. The worked example traces Dr. Holloway's toolkit across the seven components, demonstrating how the toolkit's components are acquired, maintained, and deployed over the career.

The reference works component of Dr. Holloway's toolkit is built from her university education and her professional reading. Dr. Holloway's Ph.D. in hydrology from the University of Portsmith (2014) provided the foundational scientific literature: the Journal of Hydrology, the Hydrology and Earth System Sciences, the Water Resources Research, the hydrology textbooks (the "Applied Hydrology" by Chow, Maidment, and Mays, the "Hydrology: An Introduction" by Brutsaert, the "Handbook of Hydrology" by Maidment). Her professional reading, supported by the office's training budget, includes the Remote Sensing of Environment, the ISPRS Journal of Photogrammetry and Remote Sensing, and the International Journal of Applied Earth Observation and Geoinformation. Her technical documentation includes the HEC-HMS and the HEC-RAS documentation (the Hydrologic Engineering Center's user manuals and applications guides), the WOFOST documentation, and the ENVI documentation. Her standards include the ISO 19100 series, the OGC standards, and the Federal Geographic Data Committee standards (the FGDC content standards for the digital geospatial metadata). Her professional publications include the Photogrammetric Engineering and Remote Sensing (the ASPRS journal), the ISPRS Newsletter, and the OGC Newsletter.

The software component of Dr. Holloway's toolkit is built from her university training, the submodule courses, and the on-the-job training. Her GIS software includes the ESRI ArcGIS Pro and the QGIS, with the proficiency acquired through the spatial data infrastructure course (treated in Chapter 31, section 31.3) and the daily use. Her remote sensing software includes the ENVI and the

Sentinel Toolbox, with the proficiency acquired through the processing engine course and the daily use. Her statistical analysis software includes the R and the Python, with the proficiency acquired through her Ph.D. training and the advanced course on the machine learning for the crop yield forecasting (treated in Chapter 31, section 31.9, as part of Mr. Adekoya's continuing education plan, but also available to Dr. Holloway). Her domain-specific software includes the HEC-HMS and the HEC-RAS, with the proficiency acquired through the hydrology course (treated in Chapter 31, section 31.4) and the advanced HEC-RAS 2D training (2019, treated in Chapter 30, section 30.9). Her cartographic production software includes the ArcGIS Pro and the Adobe Illustrator, with the proficiency acquired through the product studio course and the daily use. Her office productivity software includes the Microsoft Word, Excel, PowerPoint, and Outlook, with the proficiency acquired through the daily use.

The data sources component of Dr. Holloway's toolkit is built from her office's data-sharing arrangements and her familiarity with the data landscape. Her satellite imagery familiarity includes the Landsat 8 and 9 (the free-and-open, for the standing flood monitoring), the Sentinel-1 (the radar, for the cloud-penetrating flood mapping, used during the March 2024 flood event, treated in Chapter 34, section 34.9), the Sentinel-2 (the optical, for the land cover and the change detection), and the Maxar WorldView-3 (the high-resolution commercial, for the infrastructure monitoring). Her aerial imagery familiarity includes the national LiDAR DEM (the 2-metre, 0.12-metre RMSE DEM, treated in Chapter 19, section 19.9). Her sensor feeds familiarity includes the rain gauge network (the rain gauge data verified during the March 2024 flood event's diagnosis, treated in Chapter 34, section 34.9), the river stage sensors (the Portsmith gauge, treated in Chapter 24, section 24.9), and the soil moisture sensors (the NASA SMAP satellite, treated in Chapter 25, section 25.9). Her base layers familiarity includes the digital elevation model, the hydrography, the administrative boundaries, and the cadastre feed (from the Riverbourne Water Authority, treated in Chapter 27). Her thematic data familiarity includes the infrastructure asset registry (from the energy regulator, treated in Chapter 26), the reservoir data (from the RWA, treated in Chapter 27, section 27.9), and the census data (from the national statistics agency, treated in Chapter 28).

The professional networks component of Dr. Holloway's toolkit is built from her office's inter-agency work and her professional engagement. Her internal network includes the office manager (her mentor and the leader), the technical lead (her technical reference), the resource economics specialist Mr. Adekoya (her peer, with the 2023 lateral move treated in Chapter 30, section 30.9), and the operators (her operational colleagues). Her inter-agency network includes the emergency management authority's emergency operations centre (treated in Chapter 24), the water authority's operations director (treated in Chapter 27, section 27.9), the sectoral regulator's compliance director (treated in Chapter 26, section 26.9), and the national statistics agency's methodological director (treated in Chapter 28, section 28.9). Her national professional network includes the ASPRS (the Certified Mapping Scientist certification, obtained in 2021, treated in Chapter 30, section 30.9), the American Institute of Hydrology (the Certified Hydrologist certification, obtained in 2017), and the IEEE Geoscience and Remote Sensing Society. Her international professional network includes the ISPRS (the conference attendance, the publication), the IAHS (the International Association of Hydrological Sciences, the committee membership, treated in Chapter 30, section 30.9), and the GEO (the Group on Earth Observations, through the office's participation in the GEO's activities).

The standards bodies component of Dr. Holloway's toolkit is built from her standards' consultation and her standards' participation. Her standards' consultation includes the ISO 19100 series (the metadata, the services, the GML), the OGC standards (the WMS, the WFS, the WCS), the FGDC standards (the metadata, the spatial data accuracy), and the HEC standards (the HEC-HMS and the HEC-RAS conventions). Her standards' participation includes the OGC Hydrology Domain Working Group (the working group that develops the OGC standards for the hydrological data, with Dr. Holloway's participation as the office's representative), and the IAHS Hydrological Sciences division (the division that develops the hydrological standards and the guidelines, with Dr. Holloway's participation as the committee member).

The continuing education component of Dr. Holloway's toolkit is built from her ongoing training and her self-directed learning. Her continuing education includes the HEC-RAS 2D advanced training (2019, treated in Chapter 30, section 30.9), the EFAS training (2020), the ASPRS CMS certification (2021), the IAHS conference (2022, with the paper presentation), the Sentinel-2 L2A update course (February 2026, as part of the office's continuing education plan), the machine learning advanced course (May 2026), and the ESA Living Planet Symposium (June 2026, with the paper presentation). Her self-directed learning includes the professional reading (the journals, the conference proceedings), the online courses (the Coursera, the edX), and the professional engagement (the OGC working group, the IAHS committee). Her continuing education hours accumulate at the 40-hour-per-year rate, and her certifications (the HEC, the ASPRS CMS, the AIH Certified Hydrologist) are maintained through the continuing education.

The professional development path component of Dr. Holloway's toolkit is built from her career progression and her lateral moves. Her career progression includes the hazard specialist appointment (June 2018, treated in Chapter 30, section 30.9), the senior operational rank promotion (2022), and the technical lead promotion (2026, treated in Chapter 30, section 30.9). Her lateral move includes the temporary assignment to the resource economics specialization (March-May 2023, treated in Chapter 30, section 30.9, with the assignment providing the cross-domain experience). Her leadership training includes the technical leadership training (2025, treated in Chapter 30, section 30.9) and the system configuration and validation qualifications (2024-2025). Her professional development path is the framework for her toolkit's development, and the path's progression is the mechanism by which her toolkit has been built and will continue to be built.

The toolkit's deployment during the Riverbourne Basin flood event (March 2024, treated in Chapters 24, 27, 28, 29, 32, 34) is the test of the toolkit's adequacy. Dr. Holloway's reference works provided the HEC-RAS documentation that she consulted for the model's configuration (the storm surge boundary condition, the levee failure scenarios). Her software provided the HEC-HMS and the HEC-RAS that she ran for the event-specific modelling (the 90-minute run, treated in Chapter 29, section 29.9). Her data sources provided the satellite imagery (the Sentinel-1 radar, for the inundation mapping), the sensor feeds (the rain gauges, the river stage sensors), and the base layers (the LiDAR DEM, the hydrography). Her professional networks provided the inter-agency coordination (the water authority's operations director, the emergency management authority's emergency operations centre). Her standards bodies provided the standards (the OGC standards for the data sharing, the HEC standards for the modelling). Her continuing education provided the HEC-RAS 2D advanced training (2019) that enabled her to configure the 2D model for the event. Her professional

development path provided the experience (the 4 years of the hydrological modelling, the 6 years of the office's work) that enabled her to exercise the judgment (treated in Chapter 34, section 34.9). The toolkit's deployment was the test, and the test's outcome (the effective event-driven response, the accurate diagnosis, the appropriate communication) is the demonstration of the toolkit's adequacy.

This worked example illustrates the toolkit's full operation: the reference works (the scientific literature, the technical documentation, the standards, the professional publications), the software (the GIS, the remote sensing, the statistical analysis, the cartographic production, the office productivity), the data sources (the satellite imagery, the aerial imagery, the sensor feeds, the base layers, the thematic data), the professional networks (the internal, the inter-agency, the national, the international), the standards bodies (the international, the national, the professional, with the participation), the continuing education (the refresher, the update, the advanced, the self-directed), and the professional development path (the career progression, the lateral move, the leadership training). The toolkit's seven components were acquired progressively over Dr. Holloway's career (from the 2014 Ph.D. to the 2026 technical lead promotion), maintained through the continuing education and the active participation, and deployed during the Riverbourne Basin flood event (March 2024). The toolkit is the professional foundation of Dr. Holloway's career, and the foundation is the condition of her competence, her effectiveness, and her contribution to the office's geospatial intelligence. The toolkit is the professional foundation of every geospatial officer's career, and the foundation is the condition of the geospatial profession's permanence in the government. This chapter, the final chapter of the Handbook, provides the toolkit's description, and the description is the culmination of the Handbook's work. The Handbook has provided the doctrine (Parts I through VII, 36 chapters), and the toolkit (this chapter) is the resource that the officer draws on for the doctrine's application. The doctrine and the toolkit together are the geospatial officer's professional foundation, and the foundation is the condition of the officer's competence, the office's effectiveness, and the profession's contribution to the government's geospatial intelligence. The Handbook is complete, and the completion is the culmination of the work, and the work is the contribution to the geospatial profession's collective knowledge, and the contribution is the reason the Handbook exists.

Appendix A — Classification and Retention Quick Reference

This appendix provides a desk-ready quick reference for the classification regimes and the retention schedules treated throughout the Handbook. The classification regimes are organized by domain and by product type, and the retention schedules are organized by product type. The quick reference is intended for the officers who need to apply the classification and the retention in the daily work, and the quick reference is the summary that the chapters treat in detail.

A.1 The Classification Levels

The Handbook operates under a four-level classification regime, with the levels determined by the product's content, the product's consumer, and the product's potential to compromise the public safety, the critical infrastructure, or the commercial confidentiality. The four levels, from the least to the most restrictive, are: Public, Internal, Confidential, and Security Confidential.

The Public classification applies to the products intended for the public awareness and the general distribution. The Public products do not include the operational detail, the security-sensitive information, or the commercially-sensitive information. The Public products may be published on the office's website, distributed through the media, and used in the public education. The Public classification is the default for the crop condition summaries, the general flood hazard maps, the public summaries of the post-event assessments, and the public-science minerals intelligence.

The Internal classification applies to the products intended for the government's operational use, with the operational detail that the public versions exclude. The Internal products are restricted to the authorized government personnel and the authorized operational partners. The Internal products are delivered through the secure distribution channels with the access control. The Internal classification is the default for the operational flood inundation maps, the detailed exposure reports, the operational evacuation guidance, the procurement advisories, the detailed yield forecasts, and the post-season assessments.

The Confidential classification applies to the commercial minerals intelligence, with the customer-specific information that could compromise the customer's competitive position. The Confidential products are restricted to the customer's authorized personnel and the office's authorized personnel. The Confidential products are delivered through the commercial distribution channel with the access control. The Confidential classification is the default for the commercial prospectivity maps and the commercial resource estimates.

The Security Confidential classification applies to the products that include the critical infrastructure's geolocation or vulnerabilities, the security-sensitive information whose disclosure could inform the attacks on the infrastructure. The Security Confidential products are restricted to the security-cleared personnel in the government, the operational partners, and the intelligence agencies. The Security Confidential classification is the default for the energy infrastructure risk maps, the encroachment alerts, the maintenance advisories, the asset registries, and the event-specific intelligence that affects the critical infrastructure.

A.2 The Classification by Domain

The flood risk domain (Chapter 19) operates under a tiered classification regime. The general flood hazard maps are Public. The operational flood inundation maps, the detailed exposure reports, and the operational evacuation guidance are Internal. The event-specific intelligence that affects the critical infrastructure (the substations, the water-treatment plants, the hazardous-materials sites) is Security Confidential.

The food security domain (Chapter 20) operates under a split classification regime. The weekly crop condition bulletins are Public for the general version and Internal for the operational version. The seasonal yield forecasts are Internal for the operational version and Public for the summary version. The procurement advisories are Internal, because the premature disclosure could move the market prices.

The energy infrastructure domain (Chapter 21) operates under a Security Confidential classification by default, because the geolocation of the critical energy infrastructure is security-sensitive. The infrastructure risk maps, the encroachment alerts, and the maintenance advisories are Security Confidential. The asset registry is Security Confidential.

The minerals domain (Chapter 22) operates under a mixed classification regime. The public-science minerals intelligence (the geological maps, the mineral resource assessments, the regional prospectivity studies) is Public. The commercial minerals intelligence (the customer-specific prospectivity maps and the resource estimates) is Confidential. The strategic stockpile advisory is Internal.

A.3 The Retention Schedule

The retention schedule specifies the minimum period the product must be retained and the disposition at the end of the retention period. The retention schedule is aligned with the government's records management policy and the audit authority's requirements. The retention periods are organized by the product type.

The flood inundation maps (the event-driven products) are retained for 10 years, with the permanent retention for the major events (the events exceeding the 50-year return period). The flood hazard maps (the standing products) are retained permanently, as the historical record of the flood hazard assessment. The exposure reports are retained for 10 years. The evacuation guidance is retained for 5 years.

The weekly crop condition bulletins are retained for 10 years, with the seasonal summaries retained permanently. The seasonal yield forecasts are retained for 10 years. The procurement advisories are retained for 10 years, aligned to the audit cycle of the reserve authority's procurement decisions. The post-season assessments are retained permanently.

The energy infrastructure risk maps are retained for 7 years, aligned to the sectoral regulator's compliance cycle. The encroachment alerts are retained for 7 years. The maintenance advisories are retained for 7 years. The asset registry is retained permanently, as the authoritative record of the infrastructure.

The commercial minerals commissions are retained for 10 years after the completion, for the contractual disputes. The public-science minerals intelligence is retained permanently. The strategic stockpile advisories are retained for 10 years.

The post-incident reviews are retained permanently, as the institutional memory of the office's event response. The annual procedure reviews are retained for 10 years. The improvement records are retained permanently. The training records are retained for the officer's career plus 5 years.

The disposition at the end of the retention period is one of the three options: the permanent retention (the product is retained indefinitely, as the historical record), the archival transfer (the product is transferred to the government's archives), or the destruction (the product is destroyed in accordance with the government's disposal procedures, with the destruction documented in the disposal log). The disposition is specified by the product's type and the office's retention schedule, and the disposition is documented in the catalog and archive submodule.

Appendix B — Satellite Sensor Specifications Quick Reference

This appendix provides a desk-ready quick reference for the satellite sensors treated throughout the Handbook, with the sensors' characteristics organized by the resolution, the revisit, the spectral bands, and the data policy. The quick reference is intended for the officers who need to select the appropriate sensor for the analysis, and the quick reference is the summary that Chapter 7 (the acquisition) and Chapter 8 (the base layers) treat in detail.

B.1 The Free-and-Open Optical Sensors

The Landsat 8 and Landsat 9 (the NASA and USGS satellites, treated in Chapter 7, section 7.2) provide the 30-metre multispectral and the 15-metre panchromatic imagery, with the 16-day revisit (the 8-day combined revisit with the two satellites in the same orbit). The Landsat satellites carry the Operational Land Imager (the 9-band multispectral, including the visible, the near-infrared, the shortwave-infrared, and the coastal/aerosol band) and the Thermal Infrared Sensor (the 2-band thermal). The Landsat data is free and open, distributed through the USGS EarthExplorer and the Google Earth Engine. The Landsat archive, beginning in 1972 with the Landsat 1, is the longest continuous satellite imagery archive in the world.

The Sentinel-2A and Sentinel-2B (the ESA Copernicus satellites, treated in Chapter 7, section 7.2) provide the 10-metre, 20-metre, and 60-metre multispectral imagery, with the 5-day revisit (the 10-day single-satellite revisit, the 5-day combined revisit with the two satellites in the same orbit). The Sentinel-2 satellites carry the Multispectral Instrument (the 13-band multispectral, including the visible, the near-infrared, the red-edge, the shortwave-infrared). The Sentinel-2 data is free and open, distributed through the Copernicus Open Access Hub and the Copernicus Data and Information Access Services. The Sentinel-2 archive begins in 2015 with the Sentinel-2A.

The MODIS (the Moderate Resolution Imaging Spectroradiometer, on the NASA Terra and Aqua satellites, treated in Chapter 7, section 7.2) provides the 250-metre, 500-metre, and 1-kilometre multispectral imagery, with the daily revisit (the 2 daily acquisitions, one from the Terra and one from the Aqua). The MODIS carries the 36-band multispectral, including the visible, the near-infrared, the shortwave-infrared, and the thermal. The MODIS data is free and open, distributed through the NASA LP DAAC. The MODIS archive begins in 2000 with the Terra MODIS.

B.2 The Free-and-Open Radar Sensors

The Sentinel-1A and Sentinel-1B (the ESA Copernicus satellites, treated in Chapter 7, section 7.2) provide the C-band Synthetic Aperture Radar (SAR) imagery, with the 6-day revisit (the 12-day single-satellite revisit, the 6-day combined revisit with the two satellites in the same orbit). The Sentinel-1 satellites carry the C-band SAR (the 5.405 GHz, the 5.6 GHz centre frequency), with the Interferometric Wide swath, the Extra Wide swath, and the Wave modes. The Sentinel-1 data is free and open, distributed through the Copernicus Open Access Hub. The Sentinel-1 archive begins in 2014 with the Sentinel-1A. The Sentinel-1's radar imagery penetrates the cloud cover, making it the primary sensor for the flood mapping in the cloudy conditions.

B.3 The Commercial High-Resolution Optical Sensors

The Maxar WorldView-3 (treated in Chapter 21, section 21.3) provides the 0.31-metre panchromatic and the 1.24-metre multispectral imagery, with the 1-day revisit (at the high latitude) or the 4.5-day revisit (at the equator). The WorldView-3 carries the 8-band multispectral (the visible, the near-infrared, the red-edge, the coastal, the yellow), the 8-band shortwave-infrared, and the 12-band CAVIS (the Clouds, Aerosols, Vapors, Ice, and Snow). The WorldView-3 is the primary sensor for the energy infrastructure monitoring (the 0.31-metre resolution is sufficient for the sub-metre encroachment detection). The WorldView-3 data is commercial, priced per square kilometre, with the tasking through the Maxar API.

The Airbus Pleiades Neo (treated in Chapter 21, section 21.3) provides the 0.30-metre panchromatic and the 1.20-metre multispectral imagery, with the daily revisit (with the two satellites in the same orbit). The Pleiades Neo carries the 6-band multispectral (the visible, the near-infrared, the red-edge, the deep blue). The Pleiodes Neo data is commercial, priced per square kilometre, with the tasking through the Airbus API.

The Planet PlanetScope (treated in Chapter 21, section 21.8) provides the 3-metre multispectral imagery, with the daily revisit (the PlanetScope constellation of approximately 200 satellites provides the daily global coverage). The PlanetScope carries the 4-band multispectral (the visible, the near-infrared) or the 5-band multispectral (with the red-edge). The PlanetScope data is commercial, with the subscription-based access. The PlanetScope's daily revisit makes it the primary sensor for the near-real-time vegetation monitoring.

B.4 The Sensor Selection Guide

The sensor selection is the judgment that the officer exercises based on the analysis's requirements. The selection criteria comprise the spatial resolution (the analysis's scale), the temporal resolution (the analysis's frequency), the spectral resolution (the analysis's band requirements), and the data policy (the free-and-open or the commercial).

For the flood inundation mapping (Chapter 19), the primary sensor is the Sentinel-1 (the radar, for the cloud-penetrating capability), with the Sentinel-2 (the optical, when the cloud cover permits) as the secondary. The high-resolution commercial imagery (the Maxar WorldView-3, the Airbus Pleiades Neo) is used for the detailed damage assessment after the event.

For the crop condition assessment (Chapter 20), the primary sensor is the Sentinel-2 (the 10-metre multispectral, with the red-edge bands for the vegetation analysis), with the Landsat 8/9 (the 30-metre multispectral) as the secondary for the historical comparison. The Planet PlanetScope (the 3-metre, the daily revisit) is used for the near-real-time monitoring.

For the energy infrastructure monitoring (Chapter 21), the primary sensor is the Maxar WorldView-3 (the 0.31-metre, for the sub-metre encroachment detection), with the Airbus Pleiades Neo (the 0.30-metre) as the alternative. The Sentinel-1 (the radar, for the ground subsidence monitoring through the InSAR) is the secondary.

For the minerals prospectivity (Chapter 22), the primary sensors are the Landsat 8/9 (the 30-metre multispectral, for the regional alteration mapping) and the ASTER (the 15-metre multispectral,

with the shortwave-infrared bands for the mineral mapping). The airborne geophysics (the aeromagnetic, the radiometric) is the complementary data source.

Appendix C — Standards and Specifications Directory

This appendix provides a desk-ready directory of the international and national standards and specifications referenced throughout the Handbook. The directory is organized by the standard type (the geographic information, the remote sensing, the domain-specific, the reporting, the government style), and the directory is the summary that the chapters reference in detail.

C.1 The Geographic Information Standards

The ISO 19100 series (the ISO/TC 211 geographic information standards) comprises the international standards for the geographic information's structure, services, and quality. The key standards include: the ISO 19115 (the metadata, defining the metadata elements and the metadata schema), the ISO 19119 (the services, defining the service types and the service interfaces), the ISO 19136 (the Geography Markup Language, the XML grammar for the geographic information), the ISO 19142 (the Web Feature Service, the standard for the vector data's web service), the ISO 19128 (the Web Map Service, the standard for the rendered maps' web service), the ISO 19129 (the imagery, gridded and coverage data), the ISO 19157 (the data quality), and the ISO 19152 (the Land Administration Domain Model, the standard for the cadastre and the land administration).

The Open Geospatial Consortium (OGC) standards comprise the open standards for the geospatial web services. The key standards include: the Web Map Service (WMS, the OGC standard for the rendered maps), the Web Feature Service (WFS, the OGC standard for the vector data), the Web Coverage Service (WCS, the OGC standard for the raster data), the Catalog Service for the Web (CSW, the OGC standard for the catalog), the Sensor Observation Service (SOS, the OGC standard for the sensor data), and the SensorThings API (the OGC standard for the IoT sensor data).

The national standards include the Federal Geographic Data Committee (FGDC) standards in the United States (the FGDC Content Standard for Digital Geospatial Metadata, the FGDC National Standard for Spatial Data Accuracy), the Intergovernmental Committee on Surveying and Mapping (ICSM) standards in Australia (the ICSM Geographic Information standards), and the Ordnance Survey standards in the United Kingdom (the OS MasterMap standards).

C.2 The Remote Sensing Standards

The remote sensing standards comprise the standards for the satellite imagery's processing, accuracy, and quality. The key standards include: the Committee on Earth Observation Satellites (CEOS) standards (the CEOS Measurement Quality Assurance, the CEOS Calibration and Validation), the International Organization for Standardization's ISO 19129 (the imagery, gridded and coverage data), and the American Society for Photogrammetry and Remote Sensing (ASPRS) standards (the ASPRS Accuracy Standards for Digital Geospatial Data, 2014, defining the horizontal and vertical accuracy classes).

The satellite data processing levels are standardized by the CEOS and the NASA: the Level 0 (the raw data), the Level 1 (the radiometrically corrected, the georeferenced), the Level 2 (the atmospherically corrected, the surface reflectance), and the Level 3 (the mosaicked, the composited). The Landsat Collection 2 (the USGS's 2020 reprocessing) and the Sentinel-2 Level-2A (the ESA's surface reflectance) are the standard processing levels for the free-and-open optical imagery.

C.3 The Domain-Specific Standards

The hydrology standards comprise the standards for the hydrological and hydraulic modelling. The key standards include: the U.S. Army Corps of Engineers' Hydrologic Engineering Center standards (the HEC-HMS and the HEC-RAS, with the user manuals, the applications guides, and the conventions), the U.S. Federal Emergency Management Agency's flood hazard standards (the FEMA Guidelines and Specifications for Flood Hazard Mapping), and the European Commission's European Flood Awareness System (EFAS) standards.

The crop modelling standards comprise the standards for the crop growth modelling and the yield forecasting. The key standards include: the World Food Studies (WOFOST) model documentation (the Wageningen University), the Decision Support System for Agrotechnology Transfer (DSSAT) documentation (the University of Georgia), and the Food and Agriculture Organization's (FAO) crop standards (the FAO Crop Calendar, the FAO agro-ecological zones).

The mineral reporting standards comprise the standards for the mineral resource and reserve reporting. The key standards include: the Committee for Mineral Reserves International Reporting Standards (CRIRSCO) framework, the Australasian Code for Reporting of Exploration Results, Mineral Resources and Ore Reserves (the JORC Code, 2012 edition), the Canadian Institute of Mining, Metallurgy and Petroleum Definition Standards (the CIM Standards, 2014 edition), and the South African Code for the Reporting of Exploration Results, Mineral Resources and Mineral Reserves (the SAMREC Code, 2016 edition).

C.4 The Government Style Standards

The government style standards comprise the standards for the government documents' drafting and publication. The key standards include: the United States Government Publishing Office Style Manual (the GPO Style Manual, first published in 1894, with the 2016 edition as the current revision), the Australian Government Style Manual (the Style Manual, first published in 1966, with the 2020 digital edition as the current revision), and the European Commission's English Style Guide (the Style Guide, first published in 1993, with the 2022 edition as the current revision).

The plain language standards comprise the standards for the government documents' plain language. The key standards include: the United States Plain Writing Act of 2010 (requiring the federal agencies to use the plain language in the documents intended for the public), the United Kingdom's Government Communication Service's plain language guidance, and the international plain language movement (the Plain Language Association InterNational, the Center for Plain Language).

Appendix D — Comparative Institutions Directory

This appendix provides a desk-ready directory of the comparative institutions referenced throughout the Handbook, with the institutions organized by the sector and the jurisdiction. The directory is the summary that Chapter 33 (the comparative traditions) and the comparative practice sections of the domain and interface chapters treat in detail.

D.1 The National Geospatial Agencies

The United States Geological Survey (USGS), established in 1879, Reston, Virginia, with the Earth Resources Observation and Science (EROS) Center in Sioux Falls, South Dakota. The USGS is the exemplar of the minerals-funded public geoscience model and the revenue-neutral geospatial office (Chapter 33, section 33.2).

The European Space Agency's Copernicus Programme, established in 2014, with the headquarters in Brussels, Belgium. The Copernicus is the exemplar of the free-and-open data model (Chapter 33, section 33.3).

The French Institut national de l'information géographique et forestière (IGN), established in its current form in 2012, Saint-Mandé, France. The IGN is the exemplar of the cadastre-integrated mapping model (Chapter 33, section 33.4).

Geoscience Australia (GA), established in its current form in 2001, Canberra, Australia. The GA is the exemplar of the cost-recovery model in a jurisdiction with a substantial mineral sector (Chapter 33, section 33.5).

The Indonesian Badan Informasi Geospasial (BIG), established in its current form in 2011, Cibinong, West Java, Indonesia. The BIG is the exemplar of the flood-prone archipelagic state model (Chapter 33, section 33.6).

The Canada Centre for Mapping and Earth Observation (CCMEO), under Natural Resources Canada, Ottawa, Canada. The CCMEO is the exemplar of the national spatial data infrastructure model (Chapter 33, section 33.7).

D.2 The Emergency Management Institutions

The United States Federal Emergency Management Agency (FEMA), with the National Incident Management System (NIMS) and the Incident Command System (ICS), is the exemplar of the unified command model (Chapter 24, section 24.8).

The European Union's Copernicus Emergency Management Service (CEMS), operated by the European Commission's Joint Research Centre (JRC), is the exemplar of the coordinated agency model (Chapter 24, section 24.8).

Geoscience Australia's relationship with the Emergency Management Australia (EMA), within the Department of Home Affairs, is the exemplar of the embedded liaison model (Chapter 24, section 24.8).

The Japanese Cabinet Office's Disaster Management Bureau, with the Geospatial Information Authority of Japan (GSI) and the Japan Meteorological Agency (JMA), is the exemplar of the coordinated agency model with the Cabinet-level coordination (Chapter 24, section 24.8).

D.3 The Water Management Institutions

The Dutch Rijkswaterstaat, established in 1798, is the exemplar of the integrated water management model (Chapter 27, section 27.8). The Rijkswaterstaat operates under the Ministry of Infrastructure and Water Management, with the Deltares research institute as the independent analytical check.

The United States Army Corps of Engineers (USACE) and the United States Geological Survey (USGS) are the exemplars of the separated-authority model (Chapter 27, section 27.8). The USACE manages the water infrastructure, the USGS conducts the flood risk assessment, and the FEMA operates the flood response.

The Australian Bureau of Meteorology (BoM) and the state water authorities (the Melbourne Water, the Sydney Catchment Authority, the SunWater) are the exemplars of the hybrid model (Chapter 27, section 27.8).

D.4 The Strategic Reserve Institutions

The Food Corporation of India (FCI), established in 1965, with the Mahalanobis National Crop Forecast Centre (MNCFC), established in 2012, are the exemplars of the grain reserve interface (Chapter 25, section 25.8).

The National Food Authority (NFA) of the Philippines, established in 1972, with the Philippine Space Agency (PhilSA), is the exemplar of the grain reserve interface in the archipelagic state (Chapter 25, section 25.8).

The United States Strategic Petroleum Reserve (SPR), established in 1975, with the Department of Energy's Energy Information Administration (EIA) and the National Geospatial-Intelligence Agency (NGA), is the exemplar of the petroleum reserve interface (Chapter 25, section 25.8).

The United States National Defense Stockpile (NDS), managed by the Defense Logistics Agency (DLA), with the U.S. Geological Survey (USGS), is the exemplar of the critical minerals stockpile interface (Chapter 25, section 25.8).

D.5 The Training and Professional Institutions

The United States Foreign Service Institute (FSI), established in 1947, Arlington, Virginia, is the exemplar of the rank-progressive, role-specific training pipeline (Chapter 31, section 31.8).

The Canada School of Public Service (CSPS), established in 2004, is the exemplar of the government-wide training authority (Chapter 31, section 31.8).

The French Institut national du service public (INSP, formerly the École nationale d'administration, ENA), established in 1945 (reformed in 2022), Strasbourg, is the exemplar of the elite training model (Chapter 31, section 31.8).

The German Bundesakademie für öffentliche Verwaltung (BAköV), established in 1969, Brühl, North Rhine-Westphalia, is the exemplar of the federal training model (Chapter 31, section 31.8).

The Singapore Civil Service College (CSC), established in 2001, is the exemplar of the comprehensive training model (Chapter 31, section 31.8).

The professional associations include the American Society for Photogrammetry and Remote Sensing (ASPRS), the Surveying and Spatial Sciences Institute (SSSI, Australia), the Royal

Institution of Chartered Surveyors (RICS, United Kingdom), the International Society for Photogrammetry and Remote Sensing (ISPRS), and the International Cartographic Association (ICA) (Chapter 30, section 30.8 and Chapter 36, section 36.5).

Appendix E — Worked Example Index and Cross-Chapter Continuity

This appendix provides an index of the worked examples treated throughout the Handbook, with the examples organized by the cross-chapter narrative they belong to. The Handbook's worked examples are not isolated illustrations but connected narratives that span multiple chapters, demonstrating how the same events, the same officers, and the same intelligence products appear across the doctrine's different aspects. The index is the guide for the reader who wants to trace a narrative through the Handbook.

E.1 The Riverbourne Basin Flood Event Narrative

The Riverbourne Basin flood event is the Handbook's principal cross-chapter narrative, treated in 12 chapters. The narrative begins with the Riverbourne Basin's introduction in Chapter 1 (section 1.2) and the 100-year flood hazard map for the Portsmith metropolitan area in Chapter 19 (section 19.9). The flood event itself occurs in March 2024, triggered by the 72-hour rainfall of 215 millimetres (exceeding the 100-year, 72-hour rainfall depth of 185 millimetres).

The event's response is treated in Chapter 24 (section 24.9, the emergency management authority interface), where the office's activation, the event-specific intelligence production, the evacuation orders, and the post-event review are traced. The event's coordination with the Riverbourne Water Authority (RWA) is treated in Chapter 27 (section 27.9), where the Levee Section 23's settlement, the emergency reinforcement, and the model integration are traced. The Portsmith flood exposure assessment is treated in Chapter 28 (section 28.9), where the joint production with the national statistics agency achieves the 99-per-cent accuracy.

The event's information system operation is treated in Chapter 29 (section 29.9), where the six submodules' coordination is traced. The event's post-incident review and the continuous improvement are treated in Chapter 32 (section 32.9), where the four issues, the four lessons, and the four corrective actions are identified. The event's judgment situations are treated in Chapter 34 (section 34.9), where the five judgment situations across the event's timeline are traced. The event's toolkit deployment is treated in Chapter 36 (section 36.9), where Dr. Maren Holloway's toolkit's adequacy is tested.

The event's key facts: the peak stage of 4.2 metres above the flood stage at 04:00 on 25 March 2024; the inundation extent of approximately 18 square kilometres; the affected population of approximately 38,000; the damaged buildings of approximately 4,200; the Levee Section 23's settlement of 0.15 metres and the 25-per-cent breach probability; the storm surge underestimation of 20 per cent; the 6-hour update cycle through 23-26 March; the post-event review on 5 April 2024; the model validation on 20 April 2024.

E.2 The Northfields Drought Narrative

The Northfields drought is the Handbook's second cross-chapter narrative, treated in 5 chapters. The narrative begins with the Northfields' introduction as the jurisdiction's principal agricultural region (producing approximately 60 per cent of the jurisdiction's wheat) in Chapter 20 (section 20.1).

The drought's crop condition assessment is treated in Chapter 20 (section 20.9), where the NDVI at the 22nd percentile, the EVI at the 25th percentile, and the NDWI at the 18th percentile are reported.

The drought's procurement advisory is treated in Chapter 25 (section 25.9), where the heading forecast advisory of 22 July 2024 recommends the buy of 0.20 million tonnes. The drought's advisory drafting is treated in Chapter 35 (section 35.9), where the full text of the advisory's four sections is presented. The drought's training pipeline is treated in Chapter 31 (section 31.9), where Mr. Jonas Adekoya's on-the-job training includes the Northfields wheat yield forecasting in the 2025 season (the recovery from the 2024 drought).

The drought's key facts: the seasonal rainfall 35 per cent below normal (March-July 2024); the NDVI at the 22nd percentile; the yield forecast of 2.1 tonnes per hectare (90-per-cent confidence interval 1.8 to 2.4); the reserve stocks at 1.85 million tonnes (3.7 months of consumption, below the 4.5-month target); the buy signal of 0.20 million tonnes; the harvest report of 2.05 tonnes per hectare (within the forecast's confidence interval); the 8-per-cent price increase avoided by the pre-harvest buy.

E.3 The Highmoor-Portsmith Transmission Corridor Narrative

The Highmoor-Portsmith transmission corridor is the Handbook's third cross-chapter narrative, treated in 4 chapters. The corridor is the 220-kilometre, 500-kV transmission line from the Highmoor Power Station to the Portsmith Substation, serving approximately 320,000 customers in the metropolitan region. The corridor's monitoring is treated in Chapter 21 (section 21.9), where the monthly monitoring cycle detects 12 changes, of which 11 are confirmed (7 vegetation encroachments, 3 third-party constructions, 1 ground subsidence) and 8 are high-priority.

The corridor's compliance verification is treated in Chapter 26 (section 26.9), where the sectoral regulator's compliance verification, the office's compliance verification report, the field inspection, the enforcement decision, and the $250,000 fine are traced. The corridor's risk assessment drafting is treated in Chapter 35 (section 35.4), where the security register's execution is demonstrated.

The corridor's key facts: the 220-kilometre length; the 500-kV capacity; the 60-metre right-of-way (30 metres on each side); the 15-metre vegetation clearance zone; the 12 detected changes (11 confirmed, 1 false alarm); the 8 high-priority encroachments; the vulnerability scores from 1.0 to 4.0; the compliance order's deadlines (30 days for the vegetation clearing, 60 days for the geotechnical investigation, 30 days for the legal proceedings); the $250,000 fine.

E.4 The Officer Career Narratives

Dr. Maren Holloway, the hazard specialist, is the Handbook's principal officer narrative, treated in 3 chapters. Dr. Holloway's recruitment and career progression are treated in Chapter 30 (section 30.9), from the 2018 recruitment (Ph.D. hydrology, University of Portsmith, 2014; HEC certifications, 2015; Certified Hydrologist, 2017; 4 years at the Riverbourne Water Authority) through the senior operational rank promotion (2022) to the technical lead promotion (2026). Dr. Holloway's judgment during the Riverbourne Basin flood event is treated in Chapter 34 (section 34.9). Dr. Holloway's toolkit is treated in Chapter 36 (section 36.9).

Mr. Jonas Adekoya, the resource economics specialist, is the Handbook's second officer narrative, treated in 3 chapters. Mr. Adekoya's training pipeline is treated in Chapter 31 (section

31.9), from the 2025 foundational course through the submodule and domain courses to the 2028 recertification. Mr. Adekoya's cross-training for the surge capacity is treated in Chapter 32 (section 32.9, the corrective action from the post-incident review). Mr. Adekoya's advisory drafting is treated in Chapter 35 (section 35.9).

Appendix F — Glossary and Acronyms

This appendix provides a desk-ready glossary of the terms and the acronyms used throughout the Handbook. The glossary is organized alphabetically, with the terms defined concisely and the acronyms expanded. The glossary is the quick reference for the reader who encounters an unfamiliar term or acronym, and the glossary is the summary that the chapters treat in detail.

F.1 The Glossary

Acquisition: The pipeline capability that obtains the satellite imagery, the aerial imagery, the sensor feeds, and the other data the office requires (Chapter 7).

Archive: The pipeline capability that preserves the office's outputs and maintains the catalog that makes them retrievable (Chapter 12).

Base Layer: The maintained dataset that provides the geographic context, including the digital elevation model, the land cover, the hydrography, the administrative boundaries, and the cadastre feed (Chapter 8).

Cadastre: The legally binding map of the land parcels, their boundaries, their attributes, and their ownership, maintained by the cadastral authority (Chapter 23).

Catalog: The index of the office's products, with the records specifying the product's identifier, type, title, author, date, classification, coverage, storage location, and access permissions (Chapter 12).

Change Detection: The comparison of two acquisitions of the same area to identify the pixels or the objects that have changed (Chapter 10, section 10.4; Chapter 21).

Commercial Firewall: The doctrine that separates the commercial arm from the public-service arm, preventing the commercial arm's intelligence from being compromised by the public-service arm's policy preferences (Chapter 5).

Classification: The regime that restricts the product's access based on the product's content, with the four levels: Public, Internal, Confidential, Security Confidential (Chapter 11, section 11.3).

Compound Flood Hazard: The integration of the multiple flood hazards (the riverine flooding, the coastal flooding, the storm surge) into a single assessment (Chapter 19; Chapter 33, section 33.6).

Cost-Recovery Model: The funding model where the commercial intelligence sales fund the public programs, with the revenue contributing to the office's budget (Chapter 5; Chapter 33, section 33.5).

Data Fusion: The combination of the data from multiple sources (the optical and the radar, the satellite and the airborne, the satellite and the in-situ) to produce the fused product (Chapter 9, section 9.6).

Free-and-Open Data: The data policy where the Earth observation data is distributed freely to all users, exemplified by the Copernicus programme and the Landsat programme (Chapter 8; Chapter 33, section 33.3).

Forward-Sensor Relationship: The relationship where the office produces the forward-looking intelligence (the procurement advisory) and the authority decides, distinguished from the crop-monitoring service (Chapter 25).

Geospatial Information System (GIS): The information system that hosts the office's pipeline, comprising the six submodules (Chapter 29).

Intelligence-Response Division: The division between the office's production of the intelligence and the authority's conduct of the response (Chapter 24).

Job Breakdown Sheet (JBS): The TWI Job Breakdown Sheet, the step-by-step instruction with the important steps, the key points, and the reasons (Part III).

Modeling: The pipeline capability that hosts the domain-specific analytical models, with the swappable-heads design (Chapter 10).

Plug-in: The domain-specific component of the modeling workbench, with the four plug-ins: the hydrology, the crop model, the change detection, the prospectivity (Chapter 10).

Procurement Advisory: The office's forward-sensor product, synthesizing the yield forecast with the reserve stock levels to produce a release or buy signal (Chapter 20; Chapter 25).

Production: The pipeline capability that performs the cartographic design, the symbology application, the classification marking, and the format conversion (Chapter 11).

Processing: The pipeline capability that performs the georeferencing, the calibration, the mosaicking, the projection management, the cloud masking, and the data fusion (Chapter 9).

Provenance Chain: The unbroken record from the acquisition through the archive, linking the products to their source data, processing, modeling, and production (Chapter 12).

Revenue-Neutral Model: The funding model where the commercial minerals arm funds the public-service arm, with the firewall, the smoothing mechanism, and the audit regime (Chapter 5; Chapter 33, section 33.2).

Shared Backbone: The pipeline architecture's design where the five backbone submodules (the spatial data infrastructure, the acquisition desk, the processing engine, the product studio, the catalog and archive) are the same across the four domains (Chapter 6; Chapter 29).

Swappable Heads: The pipeline architecture's design where the modeling workbench's domain plug-ins vary by the domain, with the four plug-ins (Chapter 6; Chapter 29).

Training Within Industry (TWI): The programme developed during the Second World War, comprising the Job Instruction, the Job Methods, and the Job Relations dimensions (Part III; Chapter 32).

Vulnerability Score: The score that prioritizes the detected changes for the consumer, combining the asset's criticality, the change's proximity, and the change's type (Chapter 21, section 21.3).

F.2 The Acronyms

AER: Australian Energy Regulator. AEMO: Australian Energy Market Operator. AGSO: Australian Geological Survey Organisation (now Geoscience Australia). AUSLIG: Australian Surveying and Land Information Group (merged into Geoscience Australia).

BAköV: Bundesakademie für öffentliche Verwaltung (German Federal Academy of Public Administration). BIG: Badan Informasi Geospasial (Indonesian Geospatial Information Agency). BNPB: Badan Nasional Penanggulangan Bencana (Indonesian National Disaster Management Authority). BoM: Bureau of Meteorology (Australia). BRGM: Bureau de recherches géologiques et minières (French Geological Survey). BGS: British Geological Survey.

CCMEO: Canada Centre for Mapping and Earth Observation. CEOS: Committee on Earth Observation Satellites. CEMS: Copernicus Emergency Management Service. CIM: Canadian Institute of Mining, Metallurgy and Petroleum. Copernicus: The European Union's Earth observation programme. CRIRSCO: Committee for Mineral Reserves International Reporting Standards. CSIRO: Commonwealth Scientific and Industrial Research Organisation (Australia). CSPS: Canada School of Public Service.

DSSAT: Decision Support System for Agrotechnology Transfer. DVW: Deutsche Gesellschaft für Geodäsie, Geoinformation und Landmanagement (German geospatial society).

EIA: Energy Information Administration (United States). ENA: École nationale d'administration (French, reformed into INSP in 2022). EPA: Environmental Protection Agency (United States). ERDC: Engineer Research and Development Center (USACE). EROS: Earth Resources Observation and Science Center (USGS). ESA: European Space Agency. EFAS: European Flood Awareness System.

FAO: Food and Agriculture Organization (United Nations). FCI: Food Corporation of India. FEMA: Federal Emergency Management Agency (United States). FGDC: Federal Geographic Data Committee (United States). FIG: International Federation of Surveyors. FSI: Foreign Service Institute (United States).

GA: Geoscience Australia. GCO: GEOGLAM: Group on Earth Observations Global Agricultural Monitoring. GEO: Group on Earth Observations. GIS: Geographic Information System. GML: Geography Markup Language. GPO: Government Publishing Office (United States). GSDI: Global Spatial Data Infrastructure Association. GSI: Geospatial Information Authority of Japan.

HEC: Hydrologic Engineering Center (USACE). HEC-HMS: Hydrologic Engineering Center's Hydrologic Modeling System. HEC-RAS: Hydrologic Engineering Center's River Analysis System.

IAHS: International Association of Hydrological Sciences. ICA: International Cartographic Association. ICSM: Intergovernmental Committee on Surveying and Mapping (Australia). IGN: Institut national de l'information géographique et forestière (France). INSEE: Institut national de la statistique et des études économiques (France). INSP: Institut national du service public (France, formerly ENA). ISPRS: International Society for Photogrammetry and Remote Sensing. ICS: Incident Command System (United States).

JMA: Japan Meteorological Agency. JORC: Joint Ore Reserves Committee (the JORC Code, Australasia). JRC: Joint Research Centre (European Commission).

Landsat: The NASA and USGS satellite programme. LiDAR: Light Detection and Ranging.

MLIT: Ministry of Land, Infrastructure, Transport and Tourism (Japan). MNCFC: Mahalanobis National Crop Forecast Centre (India). MODIS: Moderate Resolution Imaging Spectroradiometer (NASA).

NDVI: Normalized Difference Vegetation Index. NDWI: Normalized Difference Water Index. NFA: National Food Authority (Philippines). NGA: National Geospatial-Intelligence Agency (United States). NIMS: National Incident Management System (United States). NIST: National Institute of Standards and Technology (United States). NMIC: National Minerals Information Center (USGS). NPPD: National Protection and Programs Directorate (United States, renamed CISA). NRCan: Natural Resources Canada.

OGC: Open Geospatial Consortium. OGC WMS: Web Map Service. OGC WFS: Web Feature Service. OGC WCS: Web Coverage Service. OGC CSW: Catalog Service for the Web.

PCII: Protected Critical Infrastructure Information (United States). PhilSA: Philippine Space Agency.

RCS: Rijkswaterstaat (Dutch national water authority). RADARSAT: The Canadian SAR satellite programme. RICS: Royal Institution of Chartered Surveyors (United Kingdom). RWA: Riverbourne Water Authority (the Handbook's fictional water authority).

SAMREC: South African Code for the Reporting of Exploration Results, Mineral Resources and Mineral Reserves. SAR: Synthetic Aperture Radar. SSSI: Surveying and Spatial Sciences Institute (Australia). SPR: Strategic Petroleum Reserve (United States). SMAP: Soil Moisture Active Passive (NASA satellite). SMED: Single-Minute Exchange of Die (Lean manufacturing).

TWI: Training Within Industry. TIN: Triangulated Irregular Network.

UAV: Unmanned Aerial Vehicle. USACE: United States Army Corps of Engineers. USGS: United States Geological Survey.

WOFOST: World Food Studies (crop growth model, Wageningen University). WCS: Web Coverage Service. WFS: Web Feature Service. WMS: Web Map Service.

Index

This index provides page references for the key terms, institutions, concepts, technical methods, satellites, standards, and narrative elements treated throughout the Handbook. Page numbers refer to the printed page numbers. Entries appearing on sixteen or more pages are followed by "passim" (Latin for "scattered throughout") after the first ten page references. Sub-entries are indented under their parent entries.

100-year flood, 4, 197, 202-205, 264, 311, 327-329, 406, 453

500-year flood, 199, 256, 260, 302, 303, 406

access control, 46, 109, 113, 114, 117, 155, 163, 246, 253, 254, passim

access logging, 51, 109, 113, 114, 155, 212, 261, 339, 344, 363

acquisition, 1, 2, 5, 8, 14, 17, 24, 30-32, passim

acquisition capability, 50, 56, 106, 116, 331, 333

acquisition desk, 30, 32, 52, 56, 59, 61-63, 67, 68, passim

Adekoya, Mr. Jonas, 370-373, 386, 425, 427, 439, 454, 455

aeromagnetic survey, 3, 47, 48, 53, 85, 166, 187, 192, 194, 234, 432, 447

Agency for the Cooperation of Energy Regulators (ACER), 292

  ACER, 292-294

Airbus, 35, 59, 60, 64, 70, 217, 300, 333, 364, 432, 446

alteration (geological), 94, 228, 234, 235, 421, 446

American Institute of Hydrology (AIH), 349, 357, 439

American Society for Photogrammetry and Remote Sensing (ASPRS), 3, 33, 36, 349, 356, 434, 448, 451

  ASPRS, 33, 36-38, 349, 356-359, 361, passim

Apache Airflow, 240, 335

ArcGIS (ESRI), 102, 107, 331, 348, 423, 430, 431, 438, 439

archive, 5, 9, 30, 32, 34, 36, 39, 40, 43, 44, passim

archive capability, 50, 51, 56, 116, 155, 331, 338

arrival time, 199, 256, 261, 303, 313, 364

audit regime, 6, 43, 390, 396, 399, 457

Australian Bureau of Statistics (ABS), 23, 75, 325

  ABS, 325, 326

Australian Energy Market Operator (AEMO), 293, 457

  AEMO, 293, 457

Australian Energy Regulator (AER), 293, 457

  AER, 293, 457

Badan Informasi Geospasial (BIG, Indonesia), 11, 203, 388, 394, 450, 457

  BIG, 11, 203, 388, 394, 395, 397-399, 450, 457

base layer, 4, 5, 19, 31, 49-52, 54, 55, passim
base layer capability, 50, 51, 56, 106, 116, 331, 332
Basisregistraties (Netherlands), 326
boundary dispute, 72, 73, 239, 246, 307, 308, 310
BRGM (France), 19, 392, 457
    BRGM, 22, 392, 435, 457
British Geological Survey (BGS), 13, 40, 95, 233, 457
    BGS, 40, 95, 232-234, 435, 457
building exposure, 257, 317, 319, 321, 327, 329
Bundesakademie für öffentliche Verwaltung (BAköV), 369, 370, 451, 457
    BAköV, 370, 451, 457
Bureau of Land Management (BLM), 12, 72, 75, 248, 294
    BLM, 294
Bureau of Meteorology (BoM, Australia), 75, 309, 451, 457
    BoM, 309, 310, 451, 457
cadastral authority, 18-20, 26, 27, 71, 72, 76-78, passim
cadastre, 4, 5, 8, 11, 15, 18-21, 23, passim
    cadastral, 5, 9, 18-20, 26, 27, 71, 72, 75, passim
Canada Centre for Mapping and Earth Observation (CCMEO), 388, 395, 450, 458
    CCMEO, 388, 395-399, 450, 458
Canada School of Public Service (CSPS), 369, 451, 458
    CSPS, 369, 370, 451, 458
Capital Depot, 278, 418, 426
career path, 31, 32, 38, 349-351, 355-357, 436
cartographic communication, 423, 424
catalog, 9, 30, 32, 43, 46, 48, 50, 52, 56, 58, passim
catalog and archive, 9, 30, 32, 43, 46, 48, 52, 56, 58, 62, passim
change detection, 27, 32, 37, 38, 52, 57, 58, 60, 62, 82, passim
CIM Standards, 421, 449
Civil Service College (CSC, Singapore), 369, 370, 451
    CSC, 370, 451
classification, 5, 6, 9, 17, 29, 31, 33, 38, 39, 42, passim
    Confidential, 5, 47, 57, 165, 167, 183, 192, 194, 202, 216, passim
    Internal, 9, 107, 129, 164, 178, 179, 199, 202, 204, 212, passim
    Public, 9, 161, 185, 199-202, 204, 212, 229, passim
    Security Confidential, 54, 57, 58, 102, 104, 110, 112-114, 159, passim
CleanSeaNet (EMSA), 93, 223
clearance zone, 216-218, 220, 223, 224, 283, 294-296, 419, 454

cloud masking, 23, 50, 54, 57, 79, 83, 84, 87, 88, 139, passim
École nationale d'administration (ENA), 369, 451, 458
ENA, 369, 370, 451, 458
Colonial Pipeline cyberattack (2021), 216
commercial charge code, 42, 45-48, 127, 129, 152, 153, 166, passim
commercial firewall, 12, 13, 17, 31, 39, 65, 118, 226, 230, 236, passim
firewall test, 404, 405, 409, 412, 414
commercial release, 14, 17, 27, 31, 33, 39, 42, 44-46, passim
Committee on Earth Observation Satellites (CEOS), 434, 448, 458
CEOS, 361, 434, 448, 458
common operating picture, 252-254, 258, 259, 262, 264-266, 321, 343, 385
compliance verification report, 287, 288, 291, 294, 295, 297, 425, 454
compound flood hazard, 11, 203, 394, 397-399, 456
continuous improvement, 331, 374, 376, 379-382, 437, 453
annual procedure review, 106, 376, 377, 379, 381-384, 400, 444
discipline of not freezing, 378, 379
improvement record, 375-379, 381, 382, 386, 444
method-improvement proposal, 375, 379-381, 383, 384
Copernicus Emergency Management Service (CEMS), 168, 222, 262, 293, 408, 450, 458
CEMS, 222, 262, 263, 293, 450, 458
Copernicus Programme, 2, 24, 26, 27, 49, 52-54, 56, 60, passim
Copernicus Data and Information Access Services (DIAS), 85, 390
Copernicus Open Access Hub, 27, 67, 70, 103, 110, 145, 153, 390, 445
cost-recovery model, 392, 393, 397, 450, 456
CRIRSCO, 421, 449, 458
critical infrastructure, 9, 94, 107, 128, 129, 167, 199, 200, 202, 203, passim
Critical Infrastructure Information Act (2002), 222, 261, 266, 291, 292, 329, 342, 361, 459
Protected Critical Infrastructure Information (PCII), 222, 292, 459
crop condition bulletin, 3, 7, 27, 38, 45, 55, 57, 58, 88, 99, passim
cross-section (channel), 69, 90, 299, 301, 302, 311
CSIRO (Australia), 293, 458
curve number, 90, 97, 98, 117, 169, 178, 197, 204, 205, 336, 364, 402
curve number method, 90, 97, 98, 178, 197, 204, 364
SCS curve number, 90, 97, 98, 117, 169, 178, 197, 204, 205
data fusion, 23, 50, 79, 84, 85, 88, 143, 317, 335, 363, 431, 456, 457
data pipeline operator, 8, 17, 29-32, 42, 46, 55, 56, passim
data-sharing arrangement, 209, 214, 259, 267-269, 305, 310, 315, 317, passim
data-sharing agreement, 254, 269, 284, 300, 317, 318, 323, 324, 386

Defense Logistics Agency (DLA), 277, 451
DLA, 277, 278, 451
Delft3D, 300, 336, 364
Deltares, 21, 309, 451
Digital Earth Australia (DEA), 326, 393
digital elevation model (DEM), 3-5, 8, 50, 52, 67, 69, 70, 89, passim
DEM, 4, 8, 69-71, 73, 76, 80, 90, 97, passim
Directive 2008/114/EC, 222, 223, 293
DSSAT, 52, 92, 208, 336, 348, 364, 372, 430, 431, 449, 458
eCognition, 3, 356, 368, 369, 394, 405, 412, 414, 431
Emergency Management Australia (EMA), 262, 450
EMA, 263, 450
emergency management authority, 4, 7, 9, 12, 16, 17, 27, 55, 64, 74, passim
encroachment alert, 7, 53, 57, 94, 167, 169, 172, 216, 219, 220, passim
encroachment detection, 27, 69, 93, 168, 216-218, 224, 296, 297, 446
Energy Information Administration (EIA), 277, 451, 458
EIA, 277, 451, 458
energy infrastructure intelligence, 5, 6, 11, 39, 112, 216, 225, 282, 291, 297
energy infrastructure, 5, 6, 11, 12, 16, 17, 38, 39, 52, 53, passim
energy sector entity, 7, 8, 27, 58, 93, 94, 96, 104, 106, 110, passim
EnhancedView programme, 60, 61, 222, 277, 389
ENTSO-E, 292-294
ENVI, 430, 431, 438
ERDAS Imagine, 430, 431
EROS Center (USGS), 13, 24, 26, 40, 233, 234, 389, 450, 458
ESRI, 2, 35, 102, 107, 331, 348, 423, 431, 438
European Flood Awareness System (EFAS), 63, 91, 168, 202, 358, 403, 407, 449, 458
EFAS, 63, 91, 358, 359, 440, 449, 458
European Maritime Safety Agency (EMSA), 93, 223
EMSA, 93, 223
European Space Agency (ESA), 2, 24, 49, 53, 60, 61, 82, 83, 85, 103, passim
ESA, 5, 54, 103, 361, 367, 372, 383, 384, 431-433, 440, 445, 448, 458
European Statistics Code of Practice, 324, 326
evacuation guidance, 9, 27, 52, 86, 104, 107, 128, 129, 167, 170, passim
evacuation route, 9, 107, 129, 179, 199, 204, 253, 255-257, 261, 264, 413, 416, 417
evacuation zone, 220, 253-255, 257, 264, 265, 329, 342, 343, 406, 412, 413, 416
event-driven production, 5, 8, 55, 63, 99, 120, 125, 167, 169, 170, passim
EVI, 3, 52, 88, 91, 193, 207, 208, 210, 214, 215, passim

exploration licence, 237, 242, 282, 293, 294, 412
exposure report, 9, 20, 26, 27, 52, 69, 70, 73, 77, 91, passim
Federal Emergency Management Agency (FEMA), 4, 168, 202, 262, 309, 403, 407, 449, 450, 458
    FEMA, 201, 202, 309, 407, 449-451, 458
Federal Energy Regulatory Commission (FERC), 292
    FERC, 292
Federal Geographic Data Committee (FGDC), 75, 239, 248, 249, 430, 438, 448, 458
    FGDC, 75, 438, 440, 448, 458
flood duration, 298, 303, 337
flood hazard map, 64, 100, 196-203, passim
flood inundation map, 1, 3, 7, 19, 52, 60, 71, 73, 77, 86, passim
flood risk intelligence, 5, 6, 12, 196, 205, 252, 256, 266, 303, 314, 384, 394, 397
    flood risk, 4-7, 9, 12, 16, 17, 19, 20, passim
flood stage, 260, 263, 265, 298, 306, 310, 312, 384, 411, 416, 453
flood-water boundary, 196, 200, 298, 310, 314
floodgate, 298-300, 303, 306, 309, 311-314
Food Corporation of India (FCI), 276, 451, 458
    FCI, 276, 277, 451, 458
food security intelligence, 5-7, 11, 112, 206, 272, 275, 276, 281
    food security, 5-7, 11, 16, 17, 39, 51, 52, 54, passim
Foreign Service Institute (FSI, United States), 369, 451, 458
    FSI, 369, 370, 451, 458
forward-sensor relationship, 206, 267, 274, 276-278, 281, 456
free-and-open data, 388, 390, 391, 393, 395-397, 399, 450, 456
GDAL, 240, 332, 430, 431
generalization (cartographic), 20, 21, 99-101, 104-107, 163, 337, 423, 432
geochemical data, 3, 47, 48, 85, 94, 166, 182, 187, 227, 228, passim
GEOGLAM, 92, 212, 213, 458
Geological Survey of Canada (GSC), 293
    GSC, 293, 294
GeoPackage, 72, 102, 103, 254, 283, 318, 338, 363
geophysical data, 3, 46, 47, 85, 94, 111, 168, 182, 187, 203, passim
Géoportail (IGN), 249, 391
georeferencing, 23, 49-51, 54, 56, 79-81, 86, passim
Geoscience Australia, 10-18, 23, passim
GeoServer, 240, 332, 430, 431
Geospatial Information Authority of Japan (GSI), 263, 310, 450, 458
    GSI (Japan), 263, 310, 450, 458

geospatial information system, 56, 59, 79, 89, 99, 109, 143, 155, 331, 340, passim
GeoTIFF, 9, 27, 47, 57, 58, 102-104, 107, 108, passim
GIS Certification Institute (GISCI), 33, 36, 435
    GISCI, 33, 36, 37
    GISP certification, 33, 435
government style manual, 421-423, 449
GPO Style Manual, 422, 449
GRASS GIS, 54, 102, 431
ground control point (GCP), 79-81, 147, 335, 363
    GCP, 80
ground subsidence, 24, 57, 93, 101, 216-218, 221, 224, 294, passim
Group on Earth Observations (GEO), 34, 92, 212, 352, 434, 439, 458
    GEO, 361, 430, 434, 439, 458
growing degree days (GDD), 91, 193, 208, 336, 364, 371
    GDD, 208, 209
hazard specialist, 8, 17, 27, 29, 30, 32, 37, 38, 57, 58, passim
hazardous-materials sites, 253, 257, 261, 264, 328, 416, 443
HEC-DSS, 302, 343
HEC-HMS, 91, 97, 98, 117, 168, 178, 179, 203-205, passim
HEC-RAS, 3, 8, 9, 37, 52, 54, 90, 91, 97, 98, passim
Highmoor, 223, 294-297, 311, 313, 419, 454
Highmoor-Portsmith transmission corridor, 294, 297, 419, 454
    $250,000 fine (compliance order), 296, 454
    Highmoor Power Station, 223, 419, 454
    Portsmith Substation, 223, 419, 454
Holloway, Dr. Maren, 357-359, 384, 411, 414, 438-441, 453, 454
Homeland Security Presidential Directive 7 (HSPD-7), 94, 222, 262
    HSPD-7, 94, 222
Hurricane Harvey (2017), 216, 403
Hurricane Katrina (2005), 196, 407
hydraulic model, 3, 8, 26, 27, 52, 63, 69, 85, 89, 90, passim
hydrography, 5, 49-52, 69, 73, 75, 76, 90, passim
hydrological model, 37, 62, 69, 89-91, 97, 98, 108, 110, passim
IGN (France), 4, 10, 11, 19, 29, 34, 72, 248, 356, 388, 391, 392, 450, 458
    IGN, 11, 15, 16, 18, 19, 22, 44, 72, 101, 105, passim
image differencing, 52, 93, 218, 336, 364, 365, 431
Incident Command System (ICS), 262, 450, 458
    ICS, 262, 450, 458

independence protection, 315, 318, 323-327, 330
infrastructure risk map, 5, 7, 53, 57, 94, 99, 101, 104, 106, 110, passim
InSAR, 70, 218, 300, 367, 446
Institut national du service public (INSP), 370, 451, 458
    INSP, 370, 451, 458
intelligence-decision division, 226, 267, 278, 281
intelligence-regulation boundary, 216, 226, 282, 286, 292-294, 297
intelligence-response division, 12, 200, 216, 252, 263, 266, 457
inter-agency framework, 125, 247, 249, 252, 254, 255, 269, 283-285, passim
International Association of Hydrological Sciences (IAHS), 37, 358, 439, 458
    IAHS, 38, 358, 359, 439, 440, 458
International Cartographic Association (ICA), 434, 452, 458
    ICA, 434, 452, 458
International Energy Agency (IEA), 277
    IEA, 277
International Federation of Surveyors (FIG), 434, 435, 458
    FIG, 367, 434, 435, 458
International Society for Photogrammetry and Remote Sensing (ISPRS), 3, 33, 36, 349, 356, 434, 452, 458
    ISPRS, 33, 34, 36, 37, 349, 356, 361, 367, 369, 430, 434, 438, 439, 452, 458
inundation depth, 91, 241, 255, 327, 328
inundation extent, 8, 20, 69, 91, 96, 98, 106, 107, 111, 129, passim
ISO 19100 series, 361, 430, 435, 438, 440, 448
ISO 19115 (metadata), 254, 430, 435, 448
ISO 19136 (GML), 430, 435, 448
ISO 19152 (LADM), 238, 239, 448
Japan Meteorological Agency (JMA), 310, 450, 458
    JMA, 310, 450, 458
joint flood risk assessment, 304-306
joint production, 317-319, 434, 453
Joint Research Centre (JRC), 63, 91, 92, 168, 182, 212, 222, 262, 293, 358, 390, 450, 458
    JRC, 212, 293, 294, 361, 450, 458
JORC Code, 95, 182, 194, 227, 229, 230, 234-236, 421, 449, 458
Kadaster (Netherlands), 19, 21, 23, 72, 247, 248, 326, 435
Kaizen tradition, 380, 382-384
    Imai, Masaaki, 380
KML, 102, 103, 254, 264, 338, 343, 363, 416, 417, 423
kriging, 95, 227, 228, 336, 365, 372, 420, 422

land cover, 5, 23, 49-52, 69, 73, 75, 76, passim
Landsat, 2, 13, 24, 40, 49, 54, 56, 60, 64, 80, passim
  Landsat 1 (1972), 2, 445
lateral move, 350, 358, 359, 436, 439-441
Lean tradition, 380-384, 403, 459
  5 Whys, 380, 381
  fishbone diagram, 380
  Hansei (reflection), 383
  poka-yoke, 380
  Single-Minute Exchange of Die (SMED), 380, 459
  Toyota Motor Company, 374, 380, 383, 384
  value-stream mapping, 380, 381, 383, 384
levee, 69-71, 90, 98, 178, 196-198, 200, passim
LiDAR, 3, 5, 8, 18, 67, 69, 71, 73, 76, 97, passim
lithology, 421
logistic regression, 336, 365, 372
Mahalanobis National Crop Forecast Centre (MNCFC), 213, 276, 451, 458
  MNCFC, 276, 277, 451, 458
maintenance advisory, 7, 238
Manning's n, 3, 54, 69, 91, 97, 98, 111, 117, 132, 140, passim
Maxar, 2, 35, 54, 57, 59, 60, 62, 64, 67, 68, passim
Melbourne Water, 309, 451
MIKE (DHI), 3, 37, 52, 90, 91, 198, 300, 374, 394
mineral occurrence, 95, 96, 165, 166, 194, 229, 233, 234, 291, 337, 365, 420
minerals intelligence, 5, 6, 182, 282, 287, 288, 291-294, passim
  mineral prospectivity, 1, 3, 7, 16, 19, 30, 31, 41, 45, 46, passim
mining permit, 227, 282, 293, 294
model validation, 30, 89, 95, 118, 124, 127, 131, 133, 135, 139, passim
  joint validation, 301-303, 314
model validation report, 303-305, 307
modeling, 5, 9, 21, 26, 31, 32, 37, 38, 46, 49, passim
modeling capability, 49-53, 55, 56, 106, 116, 331
modeling workbench, 31, 32, 37, 38, 46, 50, 52, 56, 57, 63, passim
MODIS, 81, 110, 213, 389, 395, 432, 445, 458
mutual TLS authentication, 255, 285, 318, 343
National Defense Stockpile (NDS), 277, 451
  NDS, 277, 278, 451
National Food Authority (NFA, Philippines), 276, 451, 458

NFA, 276, 277, 451, 458
National Geospatial-Intelligence Agency (NGA), 1, 35, 60, 94, 222, 277, 451, 458
NGA, 1, 60, 277, 278, 292, 355, 451, 458
National Incident Management System (NIMS), 262, 309, 450, 458
NIMS, 262, 263, 450, 458
National Minerals Information Center (NMIC), 95, 233, 277, 389, 458
national statistics agency, 8, 10, 22, 299, 315, 316, 324, 325, 327, 329, passim
Natural Resources Canada (NRCan), 293, 395, 450, 458
NRCan, 395, 458
NDVI, 3, 52, 57, 81-83, 88, 91, 92, 132, passim
NDWI, 52, 57, 88, 184, 208, 210, 214, 215, 278, 279, passim
neural networks, 95, 336, 365, 372
North American Electric Reliability Corporation (NERC), 292
NERC, 40, 233, 292
Northeast Blackout (2003), 216, 225, 297
Northfields, 213-215, 278, 279, 281, 359, 371-373, passim
0.20 Mt buy recommendation, 280, 281, 418, 426, 427, 454
1.85 Mt reserve stocks, 278, 279, 281, 418, 426, 454
2.1 t/ha yield forecast (Northfields), 279, 281, 417, 418, 425, 426, 454
Northfields Depot, 278, 418, 426
object-based change detection, 52, 225, 295, 297, 364
office manager, 8, 9, 14, 17, 21, 25-27, 29, 30, passim
Open Geospatial Consortium (OGC), 36, 238, 430, 434, 435, 448, 459
OGC, 102, 238, 333, 361, 362, 369, 391, 430, 434, 435, 438, 440, 448, 459
Ordnance Survey (UK), 11, 13, 14, 21, 33, 37, 40, 41, 43, 44, passim
Orfeo ToolBox, 431
orthorectification, 80, 335, 432
paragenesis, 421
payscale modifier, 35, 354-356, 358
phenological stage, 91, 92, 193, 207-209, 272, 336, 337, 364, 372
Philippine Space Agency (PhilSA), 277, 451, 459
PhilSA, 277, 451, 459
pipeline (infrastructure), 5, 8, 17, 24, 29-32, 36, 38, passim
pipeline architecture, 5, 38, 56, 131, 143, 155, 168, 182, 331, 352, 358, 361, 371, 437, 457
plain language, 421, 422, 449
plain language standard, 421, 422, 449
Plain Writing Act (2010), 422, 449
Planet (company), 35, 59, 62, 64, 68, 223, 277, 292, 333, 367, 372, 432, 440, 446

PlanetScope (Planet), 223, 432, 446
Pleiades Neo (Airbus), 59, 60, 64, 217, 300, 364, 432, 446
plug-in (domain), 5, 8, 26, 27, 32, 33, 37, 38, 50, 52, passim
    change detection plug-in, 27, 32, 37, 38, 52, 57, 89, 93, 132, 168, passim
    crop model plug-in, 27, 32, 52, 57, 88, 91, 132, 182, 184, 194, passim
    hydrology plug-in, 8, 26, 27, 32, 37, 38, 52, 89, 132, 168, passim
    prospectivity plug-in, 32, 53, 89, 94, 132, 165, 181, 182, 187, 227, 336, 363
population exposure, 199, 203, 257, 259, 317, 319-323
Portsmith, 203-205, 223, 263-265, 278, 294, 297, passim
    Portsmith Depot, 278, 418, 426
post-event assessment, 257-261, 265, 329, 330, 344, 442
post-event review, 117, 252-254, 260, 265, 266, 284, 299, 300, passim
post-incident review, 126, 129, 175, 283, 284, 290, 296, 297, 376, 377, passim
post-season assessment, 273, 442, 443
PostGIS, 240, 332, 362, 430, 431
Presidential Policy Directive 21 (PPD-21), 222
    PPD-21, 222
processing, 5, 8, 9, 17, 18, 24, 25, 30-32, passim
processing capability, 50, 51, 56, 79, 88, 106, 116, 331, 335
processing engine, 8, 30-32, 46, 52, 56, 57, 67, 79, passim
procurement advisory, 7, 91, 92, 104, 106, 181, 183, 185, 186, 192, passim
    release or buy signal, 91, 92, 183, 185, 207-209, 212, 267, 268, passim
product studio, 30, 32, 52, 56, 99-101, 103-105, passim
production, 1, 5, 6, 8, 9, 11, 17, 18, 21, 29, passim
production capability, 50, 51, 56, 106, 116, 155, 331, 337
production operator, 9, 17, 21, 29-32, 42, 46, 47, passim
professional judgment, 401
prospectivity, 1, 3, 4, 6, 7, 16, 19, 27, 30, 31, passim
prospectivity map, 3, 6, 7, 39, 40, 43, 45, 47, 53, 86, passim
provenance chain, 9, 51, 53, 54, 65, 92, 94-97, passim
pumping station, 298-300, 303, 306, 309, 311-314
QGIS, 2, 102, 332, 423, 430, 431, 438
RADARSAT, 395, 459
radiometric calibration, 57, 140
rain gauge, 8, 56, 62, 63, 66-68, 89, 97, 116, passim
rainfall-runoff, 8, 52, 89-91, 168, 170, 197, 204, 336, 364
random forest, 53, 94, 97, 228, 336, 365, 372
rating curve, 299, 301, 302, 311

Republic of Maritimea, 398-400
Maritimean Office of Geospatial Services (MOGS), 399
reservoir, 69, 73, 298-304, 306, passim
resource economics specialist, 17, 29, 30, 32, 38, 46-48, 57, 87, passim
resource estimate, 40, 53, 95, 181-183, 187, 188, 192, 194, passim
resource estimation, 35, 94, 95, 182, 226-229, 234, 336, 365, 372
retention schedule, 112, 339, 340, 442-444
retention period, 112-114, 157, 161, 272, 286, 339, 340, 344, 363, 427, 443, 444
return period, 197, 199, 204, 253, 256, 303, 364, 416, 443
design return period, 253, 256, 303
revenue-neutral model, 13, 14, 39, 43, 45, 195, 236, 390, 395, 396, 398, 399, 457
revolving fund, 13, 17, 41-45, 115, 124, 233, 390, 393, 396, 397
right-of-way, 93, 217, 218, 223, 224, 294-296, 419
Rijkswaterstaat (Netherlands), 308, 309, 451, 459
Riverbourne Basin, 8, 27, 67, 77, 97, 107, 110, 116, 128, 140, passim
18 km² inundation (March 2024), 265, 310, 312, 313, 327, 329, 343, 385, 411, 453
185 mm (100-year rainfall depth), 204, 263, 310, 453
215 mm rainfall (March 2024 event), 263, 265, 310, 416, 453
38,000 residents affected, 264, 265, 310
4.2 m peak stage (March 2024), 263, 265, 310, 312, 313, 384, 411, 416, 453
Levee Section 23, 311-314, 342, 385, 413, 414, 453
Riverbourne Water Authority (RWA), 311, 357, 384, 386, 439, 453, 454, 459
role family, 3, 29-33, 35, 38, 119, 128, passim
Royal Institution of Chartered Surveyors (RICS), 33, 36, 349, 356, 434, 459
RICS, 33, 37, 349, 356, 361, 369, 434, 435, 452, 459
SAMREC Code, 421, 449, 459
seasonal yield forecast, 99, 206, 209, 267, 268, 273, 275, 443
sectoral regulator, 1, 10, 12, 58, 112, 118, 167, 172, 178, 199, passim
secure data interface, 9, 58, 105, 108, 116, 117, 129, 179, 212, 219, passim
security clearance, 341, 355
Sentinel Toolbox, 430, 431, 439
Sentinel-1, 2, 24, 67, 68, 84, 85, 93, 98, 108, 110, passim
Sentinel constellation, 390
Sentinel-1A, 445
Sentinel-2, 2, 8, 24, 27, 54, 56-58, 60, 65, passim
Sentinel-2A, 60, 65, 86, 87, 445
Sentinel-3, 2, 24, 110, 334, 390
Shapefile, 9, 58, 72, 102-104, 107, 108, 116, 165, passim

shared backbone, 49, 52, 53, 56, 58, 131, 143, 331, 341, 344, 457
shelter, 129, 179, 199, 204, 205, 253, 255, 257, 258, 261, passim
SMAP (soil moisture), 85, 279, 302, 417, 425, 439, 459
spatial data infrastructure, 3, 31, 51-53, 56, 67, 72, 74, 75, passim
spatial-statistical integration, 316-319, 326
    geospatial join, 319, 326-328, 330
standing cadence, 253, 257, 259, 260, 266, 268, 289, 290, 299, 306, 307
statistical confidentiality, 23, 317, 318, 320, 322, 328
    de-identification, 320, 330
Statistics Netherlands (CBS), 23, 326
    CBS, 23, 326
stock-to-consumption ratio, 209, 214, 268, 270, 271, 273-275, 278, 280, 281, 418, 426
storm surge, 11, 196, 198, 202-205, 264-266, passim
storm surge level, 204, 264, 311
Strategic Petroleum Reserve (SPR, United States), 277, 451, 459
    SPR, 277, 451, 459
strategic reserve authority, 12, 27, 88, 92, 104, 106, 118, 128, 164, 181, passim
stream gauge, 63, 64, 90, 97, 98, 302, 309
substation, 8, 9, 27, 93, 98, 107, 129, 179, 198, 199, passim
supply-demand analysis, 186, 207, 208, 215, 268-273, passim
    supply-demand balance, 208, 210, 214, 268, 270, 275, 279, 426
surface reflectance, 24, 57, 81, 82, 87, 88, 91, 148, 193, 206, passim
surge capacity, 260, 265, 385-387, 455
Surveying and Spatial Sciences Institute (SSSI), 33, 36, 349, 356, 434, 451, 459
    SSSI, 33, 36, 37, 349, 356, 357, 361, 369, 430, 434, 435, 451, 459
swappable heads, 52, 53, 331, 336, 341, 344, 345, 361, 371, 437, 457
symbology, 9, 20, 21, 50, 57, 99-103, passim
Synthetic Aperture Radar (SAR), 24, 67, 70, 84, 218, 445, 459
    SAR, 67, 84, 85, 98, 108, 116, 128, 140, 141, 178, passim
technical lead, 17, 29-32, 34, 37, 38, 46, 53, passim
third-party construction, 57, 93, 101, 216-218, 224, 225, 294-296, 364, 419, 454
tidal gauge, 198, 204, 205
Training Within Industry (TWI), 118, 374, 457, 459
    Job Breakdown Sheet (JBS), 56, 59, 66, 79, 89, 99, 109, 115, 118, 131, passim
    Job Instruction (TWI), 457
    Job Methods (TWI), 374, 375, 378-380, 382, 384, 457
    TWI, 374, 375, 380, 382, 384, 457, 459
transmission corridor, 4, 8, 9, 27, 61, 98, 107, 129, 179, 198, passim

Typhoon Hagibis (2019), 310
U.S. Army Corps of Engineers (USACE), 2, 10, 90, 309, 449, 451, 459
    USACE, 309, 451, 458, 459
U.S. Census Bureau, 22, 23, 325, 326
UN Fundamental Principles of Official Statistics, 23, 324
unit hydrograph, 97, 364
United States Geological Survey (USGS), 4, 10, 12-16, 18-20, passim
    USGS, 17, 63, 64, 70, 71, 75, 76, 85, 95, 101, passim
vegetation encroachment, 57, 93, 101, 217, 218, 223-225, 294-296, 300, 364, 419, 454
vegetation index, 3, 52, 81, 83, 84, 91, 140, 207, 208, 211, passim
Virginia Mason Medical Center, 383, 384
    Rapid Process Improvement Workshops, 383, 384
vulnerability score, 27, 57, 93, 101, 217-219, 221, 224, 225, passim
water authority, 141, 196, 197, 200, 298-303, passim
waterschappen (Dutch water boards), 309
Web Coverage Service (WCS), 333, 430, 435, 448, 459
    WCS, 362, 392, 435, 440, 448, 459
Web Feature Service (WFS), 72, 238, 333, 430, 435, 448, 459
    WFS, 72, 238, 240, 244, 362, 391, 435, 440, 448, 459
Web Map Service (WMS), 72, 333, 430, 435, 448, 459
    WMS, 72, 362, 391, 435, 440, 448, 459
weights of evidence, 336, 365, 372, 422
Westbourne, 311, 313, 372, 373, 412, 420
Westbourne Minerals Company, 372, 412
Western Depot, 278, 418, 426
WOFOST, 52, 92, 97, 208, 212-215, 279, 280, passim
WorldView-3 (Maxar), 2, 54, 57-60, 62, 67, 68, 153, passim
yield forecast, 27, 52, 55, 57, 70, 73, 83, 91, 92, 96, passim