
ENKIDU

MIOVSKI GROUP

undefined

Glossary

Key terms used throughout this methodology.

- Pain point** A distinct, evidence-backed operational challenge — a technology gap, process inefficiency, workforce constraint, or information deficiency — identified from stakeholder interviews. The framework identifies 68.
- Directorate** An ECCC organizational unit that experiences pain and contributes to capability delivery. The framework covers 71 directorates within 13 branches.
- Business capability** A departmental ability to deliver an outcome, drawn from ECCC's capability model — 60 capabilities under 12 groups, of which 48 are leaf-level (operational).
- Assessment criterion** One of 23 standardized capability-maturity dimensions, organized across four pillars — Information, Technology, Workforce, and Process.
- Rubric-4** The 1-4 deterministic severity scale (Low, Medium, High, Critical/Show-Stopper) applied consistently across the framework, with 0 reserved for "no relationship."
- Severity** The Rubric-4 score of how badly a pain point affects a directorate (heatmap).
- Impact** The Rubric-4 score of how badly a pain point degrades an assessment criterion (impacts assessment).
- Delivery involvement (weight)** The 0-4 score of how much a directorate contributes to delivering a capability (capability weights).
- Burden** An intermediate score: the accumulated severity x impact a directorate carries on a given criterion.
- BC Delivery Risk** The terminal output: a normalized 0-4.0 score of how much each capability's delivery is at risk on each criterion.
- Cohen's Kappa (kappa)** A statistical measure of inter-rater reliability, used here to validate that the dual independent scoring runs agree beyond chance.
- Dual independent runs (R1, R2)** Two separate scoring passes over the same evidence, the second conducted without reference to the first — the basis of the reliability analysis.

Abstract

The ECCC Departmental Digital Roadmap rests on a rigorous, evidence-based analytical framework that translates qualitative findings from approximately 200 stakeholder interviews into quantified, actionable intelligence. This document describes the four core data assets that underpin the roadmap, explains how each was developed, details the information they contain, articulates their individual and collective value, and maps the relational architecture through which they combine to produce the Business Capability Delivery Risk assessment.

Together, these four assets form a closed analytical chain: from raw interview evidence, through pain-point identification and scoring, to understanding how those pain points burden each directorate, how they propagate into capability criteria, and ultimately how they create quantified delivery risk for each of ECCC's departmental business capabilities. The framework's defining discipline is that every quantitative score traces back to specific interview evidence, and every

derivation step is explicit, reproducible, and defensible — the same evidence-and-provenance rigour MIOVSKI GROUP applies across its analytical work.

From roughly 200 conversations to a defensible risk score for every capability on every dimension – with a complete audit trail from score back to source testimony.

1. Introduction

The ECCC Departmental Digital Roadmap is founded on a rigorous, evidence-based analytical framework that translates qualitative findings from approximately 200 stakeholder interviews into quantified, actionable intelligence. This document describes the four core data assets that underpin the roadmap, explains how each was developed, details the information they contain, articulates their individual and collective value, and maps the complex relational architecture through which they combine to produce the Business Capability Delivery Risk assessment.

Together, these four assets form a closed analytical chain: from raw interview evidence, through pain-point identification and scoring, to understanding how those pain points burden each directorate, how they propagate into capability criteria, and ultimately how they create quantified delivery risk for each of ECCC's departmental business capabilities.

2. The Four Core Data Assets

The analytical framework comprises four interconnected datasets, each serving a distinct purpose while contributing to the whole. The following sections describe each asset in detail.

2.1 Pain Point Rubric-4 Heatmap

What it is

The foundational evidence layer of the entire analytical framework. A matrix that maps 68 identified pain points against 71 ECCC directorates (organized within 13 branches), scoring the severity with which each pain point affects each directorate on a 1-4 scale.

How it was developed

Through AI-assisted deterministic analysis of approximately 200 semi-structured stakeholder interviews across all ECCC branches. First, transcripts were systematically analyzed to identify recurring operational challenges, technology gaps, process inefficiencies, and workforce constraints — yielding 68 distinct pain points. Second, each interview was analyzed to determine which pain points were referenced by each directorate and with what severity: 1 (Low) for peripheral mentions; 2 (Medium) for measurable friction; 3 (High) for significant operational constraint; 4 (Critical) for operations rendered virtually non-functional. To ensure reliability, the assessment was conducted as two independent runs (R1 and R2); a Cohen's Kappa analysis yielded an overall unweighted kappa of 0.84 ("Almost Perfect") with a weighted quadratic kappa of 0.93.

Information content

4,828 rating cells (68 pain points x 71 directorates). Each non-zero cell is a confirmed, evidence-backed relationship scored by severity. The R1 and R2 rationale sheets provide the specific interview evidence supporting each score — full traceability from score back to source material.

Value

Answers: "Which directorates are affected by which operational challenges, and how severely?" It gives leadership a precise diagnostic of their pain landscape, enables cross-directorate pattern identification (systemic versus localized), and is the primary input for both the recommendation engine and the downstream delivery-risk calculation.

2.2 Pain Point Impacts Assessment

What it is

A matrix mapping each of the 68 pain points against 23 assessment criteria across four capability pillars: Information (5), Technology (6), Workforce (6), and Process (6). It quantifies how each pain point negatively impacts each criterion on the same 1-4 scale.

How it was developed

Each pain point was independently analyzed against all 23 criteria to determine whether a causal or consequential relationship exists, and its severity: 1 (Low) marginal degradation, 2 (Medium) meaningful but contained, 3 (High) significant degradation, 4 (Show Stopper) criterion rendered virtually ineffective. Two independent runs were conducted: R1 identified 398 impact relationships; R2 identified 366. Cohen's Kappa confirmed "Almost Perfect" agreement (kappa = 0.858 overall, binary kappa = 0.931 for impact detection). Each relationship carries a detailed rationale.

Information content

1,564 assessment cells (68 x 23) per run. R1 identified 398 non-zero relationships (5.9 per pain point); R2 identified 366 (5.4). Companion rationale sheets document the pain point, criterion category and name, impact score and level, and a written rationale. The 23 criteria span four pillars:

Information	Data Quality & Accuracy, Data Integration & Accessibility, Analytics & Insights, Information Governance, Knowledge Management.
Technology	Systems & Applications, Infrastructure & Platforms, Automation & AI/ML, Integration Architecture, Technology Innovation & Modernization, Cybersecurity & Resilience.
Workforce	Talent & Skills, Capacity & Utilization, Training & Development, Organizational Structure & Roles, Funding & Budget, Vendor & Partner Ecosystem.
Process	Process Standardization & Documentation, Process Efficiency & Optimization, Quality Management, Workflow Automation, Compliance & Risk Management, Collaboration & Handoffs.

Value

Answers: "How does each pain point affect capability maturity across standardized dimensions?" It translates operational complaints into structured capability impacts — the critical bridge between qualitative pain points and quantitative capability-risk scoring.

2.3 Directorate Capability Weights

What it is

A matrix mapping ECCC's 60 departmental business capabilities (under 12 groups) against all 71 directorates, scoring each directorate's involvement in delivering each capability on a 0-4 scale.

How it was developed

The capability taxonomy was derived from ECCC's mandate, program-alignment architecture, and operational structure. Each capability was assessed for which directorates contribute and with what intensity: 0 none, 1 peripheral/supporting, 2 moderate, 3 significant, 4 primary delivery owner. Weights were informed by mandate analysis, interview evidence, and organizational-structure documentation. Typically 3-8 directorates are involved per capability, with 1-2 carrying primary responsibility.

Information content

60 capability rows x 71 directorate columns = 4,260 potential cells. Of the 60 capabilities, 48 are leaf-level (operational) with assigned weights; the remaining 12 are grouping headers. Each non-zero cell is a confirmed directorate-to-capability delivery relationship with an intensity score.

Value

Answers: "Who delivers what?" It establishes the organizational architecture of capability delivery — the structural bridge between organizational units (which experience pain) and business capabilities (which must deliver outcomes). Without it, there is no way to translate directorate-level pain into capability-level risk.

2.4 Business Capability Delivery Risk

What it is

The terminal output of the analytical chain. It maps each of ECCC's 48 operational capabilities against the same 23 criteria, producing a risk score (0-4.0) quantifying how much each capability's delivery is at risk on each criterion.

How it was developed

Derived algorithmically from the three preceding assets using a quadratic-weighted aggregation formula (Section 3). For each capability-criterion pair, the formula aggregates the pain burdens of each directorate involved in delivering that capability, weighted quadratically by delivery involvement, and normalized to the 0-4 scale.

Information content

1,104 cells (48 operational capabilities x 23 criteria). Distribution: 399 low-risk (0.1-1.0), 341 medium (1.1-2.0), 165 high (2.1-3.0), 87 critical/show-stopper (3.1-4.0), and 112 at zero. Each score is colour-coded using the standard severity palette.

Value

Answers: "What is the risk to delivering each business capability, across every maturity dimension, given the current pain landscape?" It lets senior management identify which capabilities are most at risk, which dimensions are most broadly compromised, and where remediation would yield the greatest capability improvement. It is the direct input to roadmap prioritization and resource allocation.

3. Relational Architecture and Derivation of BC Delivery Risk

The four data assets are not independent — they form a directed analytical chain where each asset

builds upon its predecessors. Understanding these relationships is essential to interpreting the final risk scores and trusting their validity.

3.1 The three-relationship chain

Relationship 1 — Directorate Pain Severity (heatmap)

Answers "which pain points affect which directorates, and how badly?" For each directorate, the heatmap provides a severity score (1-4) for every pain point that affects it. A directorate might be affected by 15-40 pain points — a directorate-specific pain profile.

Relationship 2 — Pain Point Criterion Impact (impacts assessment)

Answers "how does each pain point degrade each capability criterion?" For each pain point, the assessment provides a score (1-4) for every criterion it affects. A pain point might impact 4-8 criteria — a pain-point-specific impact profile.

Relationship 3 — Directorate Delivery Involvement (capability weights)

Answers "which directorates deliver each capability, and to what degree?" For each capability, the weight matrix identifies 3-8 involved directorates with involvement scores (1-4) — a capability-specific delivery profile.

3.2 How the three relationships combine

The derivation proceeds through two computational stages.

Stage 1 — Directorate Criterion Burden

For each directorate, Relationships 1 and 2 are combined to produce a criterion-level burden score. The burden of directorate d on criterion k is:

$$\text{Burden}(d, k) = \sum [\text{Severity}(d, pp) \times \text{Impact}(pp, k)]$$

The summation runs over all pain points (pp) that affect directorate d and also impact criterion k . The product means a pain point that severely afflicts a directorate and deeply degrades a criterion contributes far more than one with low scores on either dimension. For example, "Legacy Infrastructure" at severity 4, impacting "Infrastructure & Platforms" at impact 4, contributes $4 \times 4 = 16$ — whereas a severity-1, impact-2 combination contributes only 2.

Stage 2 — Capability Risk with Quadratic Weighting

The burden scores are aggregated up to the capability level using Relationship 3, applying a quadratic weighting formula:

$$\text{Risk}(\text{cap}, k) = \sum_d [w_d^2 \times \text{Burden}(d, k)] / \sum_d [w_d^2]$$

Here w_d is the delivery-involvement weight of directorate d for the capability, and the summation runs over all directorates involved in delivering it. The quadratic weighting is the critical design choice: it ensures directorates with higher delivery responsibility contribute disproportionately more to the capability's risk profile.

3.3 Why quadratic weighting

The choice of quadratic (w^2) rather than linear (w) weighting is deliberate. Under linear weighting, a

directorate with weight 4 would contribute twice the risk of one with weight 2. Under quadratic weighting, it contributes four times the risk (16 versus 4). This reflects the operational reality that a primary delivery owner's problems have outsized impact compared to a peripheral contributor's.

Consider a capability delivered primarily by Directorate A (weight 4) with support from Directorate B (weight 2). If A suffers severe legacy-infrastructure problems while B does not, the capability's risk should be dominated by A's burden. Quadratic weighting ensures this: A's contribution is $16/(16+4) = 80\%$ of the total, whereas linear weighting would yield only $4/(4+2) = 67\%$. The quadratic approach better captures the concentration of delivery risk.

3.4 Normalization

The raw risk scores span a wide range (0 to approximately 116) because they accumulate multiple pain-impact products. To make them interpretable on the standard 0-4 scale, they are normalized using the 95th percentile of all non-zero raw scores as a ceiling: a raw score at or above the 95th percentile maps to 4.0, while lower scores are proportionally scaled. The 95th percentile (rather than the absolute maximum) prevents a few extreme outliers from compressing the scale for the majority. Scores are rounded to one decimal place.

4. Analytical Data Flow

The complete chain is a directed flow from evidence to risk. Approximately 200 stakeholder interviews are processed through AI-assisted deterministic analysis, Rubric-4 scoring, and dual independent runs to produce the Pain Point Rubric-4 Heatmap (68 pain points x 71 directorates, severity 1-4). That heatmap feeds two parallel assets: the Pain Point Impacts assessment (68 x 23 criteria) and the Directorate Weights matrix (48 capabilities x 71 directorates). The three then combine through quadratic aggregation —

$$\text{Risk}(\text{cap}, k) = \text{sum} [w^2 \times \text{sum}(\text{severity} \times \text{impact})] / \text{sum } w^2$$

— to yield the terminal Business Capability Delivery Risk matrix (48 capabilities x 23 criteria, normalized 0-4.0 scores).

5. Validation and Methodological Rigour

Dual independent runs

Both the Rubric-4 heatmap and the Pain Point Impacts assessment were conducted twice independently, the second run approaching the evidence without reference to the first — analogous to inter-rater reliability in qualitative research.

Cohen's Kappa validation

Formal inter-rater reliability analysis was performed on both datasets. The heatmap achieved kappa = 0.84 (Almost Perfect); the impacts assessment kappa = 0.858. Both substantially exceed the kappa ≥ 0.61 threshold considered acceptable for applied research.

Rationale traceability

Every non-zero score in both datasets is accompanied by a written rationale that traces back to specific interview evidence — a complete audit trail from final risk score to source testimony.

Evidence-based scoring

The Rubric-4 deterministic framework constrains scoring to evidence-supported judgments, eliminating subjective inflation or deflation. Pain points are only scored for directorates where interview evidence confirms their presence.

95th-percentile normalization

The normalization approach prevents outliers from distorting the scale while preserving meaningful differentiation across the full score range.

Quadratic sensitivity

The quadratic weighting was validated against linear alternatives and confirmed to better reflect operational reality, where primary delivery owners' burdens dominate capability risk.

6. Interpreting Business Capability Delivery Risk Scores

The 0-4.0 risk scores in the BC Delivery Risk matrix should be interpreted as follows.

3.1 - 4.0 • Show Stopper

The pain points affecting the responsible directorates are so severe, and their impact on this criterion so deep, that capability delivery on this dimension is fundamentally compromised. Requires urgent, structural intervention.

2.1 - 3.0 • High Risk

Significant pain burden flowing through high-involvement directorates creates material delivery risk. The capability can function but operates well below its potential on this dimension. Near-term remediation recommended.

1.1 - 2.0 • Medium Risk

Moderate pain burden creates friction in capability delivery. Operations continue but with inefficiency, workarounds, or quality degradation. Medium-term improvement warranted.

0.1 - 1.0 • Low Risk

Minor pain burden with limited impact. The criterion is broadly functional. Monitor and address opportunistically.

0.0 • No Risk

No identified pain points affecting the involved directorates impact this criterion, or no directorates with assessed pain are involved in delivering this capability.

A high risk score on a specific criterion for a specific capability does not necessarily mean that criterion is globally problematic — it means the particular combination of directorates delivering that capability carry pain burdens that converge on that criterion. Different capabilities may show different risk profiles even on the same criterion, because their delivery teams differ.

7. Strategic Application

Capability-driven prioritization

Rather than prioritizing remediation on frequency or severity alone, BC Delivery Risk enables

prioritization based on which pain points create the greatest risk to the most critical business capabilities.

Targeted investment

The criterion-level granularity reveals whether a capability's risk is concentrated in technology, workforce, process, or information — guiding the type of intervention required, not just its urgency.

Cross-directorate diagnosis

The relational chain reveals cases where multiple directorates sharing a capability each bring different pain points that compound into high delivery risk, indicating the need for coordinated rather than siloed remediation.

Maturity projection

As pain points are remediated, the model can be re-run to project how delivery-risk scores would change, enabling evidence-based forecasting of roadmap impact.

Accountability alignment

The directorate-to-capability mapping establishes clear accountability for delivery risk, enabling leadership to assign remediation responsibility to the directorates whose pain burdens contribute most to capability risk.

END OF DOCUMENT