

21 July 2026

Commerce Commission

infrastructure.regulation@comcom.govt.nz

c/o Elloise Kidd, Senior Analyst, Electricity Distribution

Subject: Submission: Draft Resilience Guidance Feedback

Dear Commission,

This submission provides feedback on the draft Resilience Principles: Decision-making framework for critical infrastructure providers (3 July 2026) and the accompanying ResOrgs report. I support the guidance initiative and offer constructive feedback drawn from operational experience in New Zealand water services and published literature on risk and uncertainty.

The guidance is a genuine step forward. Three gaps warrant attention:

- The framework does not distinguish between recoverable losses and ruin-class outcomes where the service ceases to exist. Cost-benefit analysis is the right tool for the former and the wrong tool for the latter.
- The uncertainty treatment addresses variability in known risks but does not extend to risk categories that have not yet been identified. Where the risk itself has not been conceived, the response is structural buffering, not better analysis.
- Deferred renewal is an undisclosed liability that compounds nonlinearly. The Information Disclosure regime should make it visible.

The submission draws on Taleb and Cirillo's work on epistemic uncertainty and the thickness of tails in statistical estimation (2019) and on the ResOrgs report's own analysis of the limitations of Value of Lost Load metrics for water services.

I am happy to discuss any aspect of this submission.

Connan Negus CMEngNZ

1. Ruin is not a cost

Cost-benefit analysis works when losses are recoverable. The entity absorbs the cost, learns, adjusts, and continues operating. Most resilience investments fall into this category and CBA is the correct evaluation tool for them.

CBA does not work when the loss eliminates the service. A sole-source water treatment plant destroyed with no alternative supply, no redundancy, and no recovery pathway is not a cost to be optimised. It is an absorbing barrier: once the system reaches zero, it cannot recover. The mathematical expectation of the game is irrelevant because the entity does not survive to collect the upside.

The guidance (paragraph 2.7) states that evidence quality should increase with the significance of the risk. I agree, and argue the principle does not go far enough. For ruin-class risks, the decision framework itself should change, not merely the quality of evidence within the same framework. Ruin requires a survival constraint that precedes optimisation. The entity first establishes the floor below which service cannot fall (minimum redundancy, minimum recovery pathway, minimum buffer), then applies CBA to everything above that floor.

The problem is compounded by the probability estimates used to assess these risks. The return-period analysis used throughout the guidance (1/100-year, 1/250-year, 1/500-year events in Principles 2 and 9) deserves scrutiny. New Zealand has less than 200 years of instrumental weather records. Fitting a 1/500-year return period from this dataset extrapolates beyond the data by a factor of more than two. Taleb and Cirillo (2019) demonstrate that epistemic uncertainty about tail behaviour causes real-world tails to be fatter than fitted models predict. In practical terms: the probability of extreme events is higher than the return-period number suggests, and the confidence interval around that probability is wider than the precision of the number implies. Presenting 1/500 as a usable design parameter without acknowledging the epistemic limitations of the underlying data creates false confidence.

Recommendation: Add a threshold screening principle. Before CBA commences, screen for irreversible or ruin-class outcomes. Where the consequence eliminates the service or creates irreversible public health harm, set a survival constraint (minimum redundancy, minimum recovery pathway) that precedes and is not subject to cost-benefit optimisation. Require return-period analysis to carry explicit acknowledgment of the limitations of the underlying data and distributional assumptions.

2. Fat tails and the limits of return-period analysis

The Table 3.1 indicator for Principle 11 states that techniques such as Monte Carlo simulation, real options analysis, and Dynamic Adaptive Policy Pathways are expected for 'significant investments and mature organisations.' Paragraph 2.37 recommends these techniques more broadly for deep uncertainty contexts, and Table 4.1 describes them without a maturity gate. The inconsistency matters, because Table 3.1 is what assessors will check against.

Two problems with the Table 3.1 formulation. First, the computational barrier to these tools has collapsed. Monte Carlo simulation requires a spreadsheet-competent analyst and free software. It is

no longer a specialist capability that only mature organisations can deploy. Maturity classification should not gatekeep tools that are now accessible.

Second, for risks with ruin-class consequences and fat-tailed distributions, these tools are not extras. They are minimum requirements. An entity with a single water treatment plant serving an isolated community faces a risk with potentially irreversible consequences. Requiring that entity to reach ‘mature’ status before it accounts for uncertainty is requiring it to ignore the most important feature of its risk profile until it has the organisational capacity to do the analysis properly. The risk does not wait for the organisation to mature.

The return periods throughout the guidance (1/100, 1/250, 1/500) are implicitly fitted to thin-tailed, stationary distributions. Under non-stationary climate conditions, these return periods are not fixed parameters. A 1/100-year flood today may be a 1/50-year flood in two decades. The guidance should require that return-period analysis explicitly state the distributional assumptions and the period of record underlying the estimate, and that sensitivity testing be applied to the return-period values themselves, not only to the consequences.

Recommendation: Remove the maturity gate in Table 3.1 on uncertainty analysis for ruin-class risks. Where the consequence of failure is irreversible, uncertainty analysis should be expected regardless of organisational maturity level. Require return-period estimates to carry explicit statements of the underlying data period, distributional assumptions, and sensitivity of conclusions to changes in those assumptions.

3. Redundancy is not inefficiency

The guidance defines ‘efficient’ (paragraph 1.15) as including public health, environmental performance, and service continuity over the long term. Redundancy that prevents ruin-class outcomes is a cost of delivering that definition of efficiency. It is not gold-plating.

The risk in price-quality regulation is that redundancy is scored as inefficient expenditure because it provides no value during normal operations. The standby pump that sits idle for ten years has delivered zero measurable output. When the primary pump fails during a storm and the standby keeps the community supplied, the value of that decade of idleness is realised in a single event. CBA of the standby pump against its utilisation rate would reject it. CBA of the standby pump against the consequence of its absence would approve it. The framing determines the answer.

The optioneering framework (Principle 7) presents avoid, control, transfer, and accept as parallel options to be weighed. They are not parallel. They are hierarchical. Avoid (prevent ruin) is a survival constraint. Accept (tolerate small losses that create learning and optionality) is a design choice about where to allow stress into the system. Control and transfer (manage recoverable risk) are the CBA domain. The sequence is: screen for ruin first and set the floor, then identify what small losses to accept, then optimise the middle.

This does not exempt ruin-prevention investment from the requirement to demonstrate robust decision-making. The survival constraint sets a floor; the entity must still demonstrate that the floor is proportionate and the investment above it is efficient.

Recommendation: Add explicit guidance that minimum redundancy for ruin prevention is a legitimate and expected component of efficient service delivery, consistent with paragraph 1.15, and should not be subject to efficiency challenge in price-quality review. Reframe the optioneering hierarchy to reflect the categorical difference between survival constraints (avoid), design choices about acceptable stress (accept), and optimisation variables (control, transfer).

4. Deferred renewal is undisclosed debt

Deferred renewal is worse than financial debt. Debt is visible on the balance sheet, voluntary, contractual, and has an agreed repayment schedule. Deferred renewal is invisible, involuntary from the consumer's perspective, non-contractual, and compounds nonlinearly. A pipe that needed \$50,000 renewal five years ago may now require \$200,000 emergency replacement plus the cost of service disruption and public health risk during the failure.

The asymmetry matters. The decision-maker who deferred the renewal captured the benefit: lower rates today. The future ratepayer inherits the cost: higher rates and worse service tomorrow, without having consented to the trade-off or even being informed of it. Unlike debt, there is no informed counterparty.

The New Zealand water sector investment shortfall has been estimated at \$120 to \$185 billion nationally (Pennington, 2021, cited in the ResOrgs report). This figure reflects, in significant part, the accumulated consequences of decades of underinvestment.

A potential gaming risk: if the indicator is the gap between forecast renewal and actual spend, entities may lower their renewal forecasts to close the gap on paper. The counter is to require disclosure of both the forecast methodology (including the asset condition data underpinning it) and the actual spend. An entity that understates its forecast has to explain why its renewal programme is low relative to asset condition.

Recommendation: Require water service providers to disclose, as part of the Information Disclosure Determination, the cumulative gap between renewal expenditure forecast in asset management plans and actual renewal spend, together with the asset condition data and methodology underlying the forecast. Make the deferred renewal liability visible to governance bodies, shareholders, and consumers.

5. Adaptive pathways as default for water infrastructure

Dynamic Adaptive Policy Pathways (DAPP) is included in the guidance as one of several decision-making frameworks (Principle 11, Chapter 4 Table 4.1). Table 4.1 notes that DAPP 'may be especially relevant for water and other sectors facing long asset lives and uncertain hazard trajectories.' Paragraph 2.37 recommends adaptive planning approaches for deep uncertainty contexts generally. Yet the Table 3.1 indicator for Principle 11 gates these techniques on 'significant investments and mature organisations.' For water infrastructure specifically, DAPP should be the expected approach rather than an advanced-maturity aspiration.

The case is straightforward. Water assets have long lives (50 to 100 years for pipelines and structures). The hazard environment is non-stationary (climate change alters flood frequency, drought patterns, and sea level within asset lifetimes). The sector's analytical maturity is low, meaning point-forecast planning is unreliable because the data to support confident single-scenario forecasts does not exist. And the consequence of getting it wrong is irreversible public health harm.

Under these conditions, committing to a single design scenario for a 50-year asset is not prudent planning. It is a bet on a future that cannot be predicted with the precision the commitment requires. DAPP, with explicit trigger points that signal when to shift pathways, is the honest response to deep uncertainty about long-lived assets.

Recommendation: Under Principle 11 and Chapter 3, elevate DAPP from 'available to mature providers' to 'expected approach for long-lived water infrastructure investment planning, proportionate to the scale and consequence of the investment.' This does not require sophisticated modelling. A trigger-based plan can be built from qualitative thresholds (for example, 'when flood frequency exceeds X events per decade, initiate pathway review').

6. Feedback on Chapter 3 indicators

Chapter 3, Table 3.1 is the operational heart of this guidance. It is what assessors will check against and what providers will build toward. Five observations:

6.1 Prescribe outcomes, not tools

Several indicators prescribe specific tools: GIS layers (Principle 2), quantitative assessment methodology (Principle 1), Monte Carlo and real options analysis (Principle 11). These are methods, not outcomes. An entity with sophisticated GIS and poor risk understanding scores well. An entity with deep operational knowledge and no GIS scores poorly.

Indicators should describe the understanding to be demonstrated, not the tool to be used. For example: 'Demonstrate spatially differentiated understanding of hazard impact on assets and consumers' rather than 'GIS layers showing impact.' The entity chooses the proportionate method; the assessor evaluates the quality of understanding.

6.2 Unknown risks are absent

Every indicator in Table 3.1 assumes the entity knows what its risks are. This is the known-known quadrant: risks identified, consequences understood, probabilities estimated. No indicator addresses the category of risks the entity has not yet identified.

This matters because the most consequential failures are typically in this category. Risks that were not in the register, dependencies that were not mapped, failure modes that had not been conceived. The response to unknown risks is categorically different from the response to known risks:

- **Known risks with known parameters:** manage through the standard cycle of rate, own, mitigate, monitor. CBA and ISO 31000 work here.

- **Known gaps (known question, unknown answer):** investigate. Convert to known-known by filling the data gap.
- **Genuinely unknown risks:** buffer and surface. Build structural capacity to absorb surprise (reserves, redundancy, essential spares, flexible contracts, trained surge capacity) and maintain active processes that surface what is not yet known (post-incident learning, structured challenge of existing assessments, cross-entity knowledge exchange, frontline intelligence).

A particularly common variant is knowledge that exists within the organisation but has not reached decision-makers — risks known to frontline staff but invisible to governance. The response is structural: mechanisms that make organisational boundaries permeable to risk intelligence.

The ResOrgs report recognises the concept of incomplete knowledge (Figure 12 uncertainty matrix; strategic stockpiling and buffer recommendations in Section B1.5). This recognition has not translated into an expected indicator or response in the guidance's Table 3.1. The third category — where the most consequential failures typically occur — has no indicator, no framework, and no expected response.

Buffer adequacy and buffer-trend monitoring deserve specific attention. An entity may declare a risk mitigated, but if the buffers that underpin the mitigation are eroding (staff capability declining, spares not replenished, financial reserves drawn down), the declared posture is decaying. Tracking buffer state over time is a leading indicator of resilience. Current condition assessments are lagging indicators.

Recommendation: Add an indicator: 'Demonstrate processes for detecting unknown risks, including structured challenge of existing risk assessments, post-incident learning systems, and cross-entity knowledge exchange. Demonstrate that structural buffers (staffing, financial, physical, knowledge) are monitored for adequacy and trend, and that buffer erosion triggers reassessment.'

6.3 Cost of inaction

Principle 9 requires consideration of the 'do nothing' option and expects it to be compared against alternatives with reasons documented. This is the right starting point. What is missing is a requirement to rigorously size the consequences of inaction, not merely list it as an option. Systems degrade. The cost of inaction is non-zero and compounds nonlinearly.

Recommendation: Require that the 'do nothing' option carries a rigorous consequence assessment, including the compounding cost of continued degradation over the planning horizon, with the same analytical standard applied to active intervention options.

6.4 Decision-change evidence

The indicators measure whether analysis was performed, not whether it changed anything. An entity can produce GIS layers, CBA reports, community engagement records, and criticality frameworks that satisfy every indicator without altering a single operational decision. The analysis becomes performative.

Recommendation: Add an indicator: ‘Demonstrate that resilience assessment outputs have influenced actual investment decisions, including decisions to change priorities, defer investments, or reject previously preferred options.’

6.5 Trends, not just snapshots

Annual reporting of indicator values creates a time series. But a time series is only useful if someone analyses the trend. An entity whose resilience metrics are improving is in a different position from one whose metrics are static or degrading, even if the snapshot values are identical.

Recommendation: Add a requirement that entities analyse and report trends in resilience metrics to governance, not merely snapshot values. Where trends show degradation, require explanation and remediation planning.

7. Responses to consultation questions

Q1: Is the guidance clear?

Largely yes. The proportionality language is well-judged. The gap: no distinction between recoverable and ruin-class risks (Section 1 above). A provider reading the guidance could reasonably conclude that CBA applies to all resilience decisions. It should be clear that some decisions precede CBA.

Q3: Implementation barriers, overlaps, or accountability issues?

For water services, the interaction between the Commerce Commission (economic regulation) and Taumata Arowai (drinking water quality regulation) creates a structural seam. Neither regulator owns the full resilience picture. A water entity could be quality-compliant but economically constrained from investing in resilience, or have resilience investment approved but face quality compliance failures because investment was directed elsewhere. This is a system-of-systems governance challenge. The guidance’s own Principle 3 (logical interdependencies) identifies the phenomenon. Acting on it requires cross-regulator coordination on how resilience investment is assessed and approved.

Q4: What should the Commission prioritise?

- A ruin-screening triage for evaluating resilience investment proposals, so that survival constraints are set before CBA commences.
- An epistemic classification requirement, so that entities acknowledge what they do not know and demonstrate appropriate responses for each category of uncertainty.
- Cross-sector interdependency information disclosure: require each regulated entity to disclose its critical upstream dependencies (what other infrastructure it requires to function) and critical downstream dependencies (what services or entities depend on it), in a standardised format. The Commission regulates entities across five infrastructure sectors and is uniquely positioned to build a cross-sector dependency map from its regulated portfolio.

- Business case guidance that incorporates the ruin-screening triage, epistemic classification, and adaptive pathways as standard elements.

Q7: Are any principles missing?

Yes. An overarching principle on epistemic humility: acknowledge what you do not know and build structural capacity to absorb surprise. The current principles assume a complete or completable risk inventory. They do not address the category of risks that have not been conceived. The response to genuine ignorance is not analysis (which requires knowing what to analyse) but structural buffering, redundancy, optionality, and active processes to surface what is hidden.

Q8: Are there missing frameworks or evaluation tools?

The guidance and the ResOrgs report operate almost entirely within the domain of known risks with estimable probabilities. The frameworks in Chapter 4 extend to known uncertainty (DAPP, RDM), which is valuable. The ResOrgs report's Figure 12 uncertainty matrix recognises ignorance and unawareness as dimensions of incomplete knowledge. What is absent from the guidance itself is a framework and indicator for genuinely unknown risk categories, where the risk itself has not been identified.

An epistemic state classification would address this. Risks can be classified by what is known about them:

- **Known-known:** risk identified, parameters understood. Response: manage. CBA, ISO 31000, conventional risk management.
- **Known-unknown:** risk identified, parameters uncertain. Response: investigate and analyse. DAPP, RDM, Monte Carlo, scenario testing.
- **Unknown-unknown:** risk category not yet identified. Response: buffer and surface. Structural redundancy, reserves, optionality, post-incident learning, structured challenge of assumptions.

Different epistemic categories require different decision frameworks and different response types. The guidance currently provides tools for the first two categories. The third category, where the most consequential failures typically occur, has no framework, no indicator, and no expected response in the guidance's Table 3.1. This is the single largest gap in the guidance.

Q10: Practical examples?

No practical examples to share at this time.

Q11: Iwi, hapū, and Māori perspectives?

Community resilience practices in smaller or isolated communities predate modern infrastructure. These practices were shaped by New Zealand's hazard environment over centuries and contain operational intelligence about local hazard patterns, self-provisioning, and community response that engineering analysis would need to derive from first principles if it were honest enough to look.

Chesterton's fence applies: before proposing infrastructure solutions that replace or override community resilience practices, understand why those practices exist and what knowledge they encode. The framework should recognise community resilience knowledge as design intelligence, not merely stakeholder input to be consulted.

Q12: Cross-sector vulnerabilities and interdependencies?

Electricity is the keystone dependency. Most other infrastructure ceases to function without it, and it tends to fail first in widespread hazard events. Water is second in consequence priority because public health impacts from loss of safe water supply are immediate and can be irreversible.

The chain from electricity to water (pumping, treatment, SCADA and telemetry communications) to telecommunications (remote monitoring, alarm systems) to roading (access for repair crews) is the critical interdependency path. The framework should weight assessment and disclosure of this chain as the highest priority for cross-sector resilience planning.

The system-of-systems nature of infrastructure means that no single entity can assess its own resilience in isolation. Resilience is an emergent property of the interactions between infrastructure systems, not a property of any individual system. The Commission's cross-sector regulatory scope makes it the natural coordinator for interdependency disclosure and assessment.

References

Pennington, P. (2021, June 2). New Zealand water infrastructure spend could reach \$185b by 2051 — studies. RNZ. <https://www.rnz.co.nz/news/national/443888/new-zealand-water-infrastructure-spend-could-reach-185b-by-2051-studies>

ResOrgs (2025). Infrastructure Resilience: Key principles for regulating the performance of utilities. Prepared for Commerce Commission, August 2025.

Taleb, N.N. & Cirillo, P. (2019). Branching epistemic uncertainty and thickness of tails. arXiv:1912.00277. <https://arxiv.org/abs/1912.00277>