Climate Disruption, Predictive Intelligence,

and the Architecture of Operational Survival in Financial Services

---

PUBLISHED BY
Fulcrum Digital

What This Paper Argues

Financial institutions have spent decades treating climate risk as a modelling problem. The wildfires, floods, and earthquakes of the past three years have demonstrated that it is an operational problem—one that appears in the interval between when disruption begins and when an institution can respond.

→ The evidence is direct. In October 2024, when flooding struck Valencia, Spain, 9% of ATMs in the flood zone went offline and 37 of 298 bank branches were forced to close. The Bank of Spain estimated that lenders carried €20 billion in direct loan exposure to the affected area. The flooding was not a surprise event; a DANA weather system had been tracked for days. The weakness sat in the operational chain connecting warning, assessment, and response.

→ In January 2024, a 7.5 magnitude earthquake struck Japan's Noto Peninsula on New Year's Day, when financial institutions operated on skeleton staffing and reduced systems monitoring. Thirty-six thousand households in Ishikawa and Toyama prefectures lost power. Financial services were suspended across multiple prefectures. Japan's Financial Services Agency issued emergency countermeasures and coordinated operational response, but the response had to be improvised because no institution had anticipated the specific timing and cascading infrastructure failures.

→ In Canada in 2024, insured losses from severe weather reached a record CAD $8.5 billion—nearly triple the previous year and 12 times the annual average from 2001 to 2010. In just two months that summer, four catastrophic events produced over CAD $7 billion in insured losses and a quarter of a million insurance claims, 50% more than Canadian insurers typically process in an entire year. The sector incurred underwriting losses in personal property for both 2023 and 2024.

FINANCIAL RESILIENCE NOW DEPENDS ON OPERATIONAL SENSING: THE ABILITY TO SEE WHERE DISRUPTION IS UNFOLDING, HOW SEVERE IT IS, AND HOW EXPOSED THE INSTITUTION IS WHILE THERE IS STILL TIME TO RESPOND.

What This Paper Covers

- Why climate disruption is now an operational systems problem for financial services, supported by evidence from five continents
- The three-layer architecture of predictive intelligence infrastructure and where the investment gap sits
- How AI and satellite-based catastrophe modeling are changing the speed and precision of exposure assessment
- What quantum methods are operationally contributing today, separate from the long-term commercial quantum horizon
- How six global regulators—OSFI, Bank of England, ECB, NGFS, APRA, and the Basel Committee—are converging around the same operational resilience requirements
- Why institutions building predictive operational infrastructure now are likely to carry compounding structural advantages as climate-driven disruption intensifies

$320B

Global economic losses from natural disasters, 2024 (Munich Re)

$223B

Global uninsured losses from weather-related disasters, 2024 (Aon via World Economic Forum)

CAD $8.5B

Canada's costliest insured severe weather year, 2024 (IBC)

15–30%

Estimated global GDP loss under NGFS climate scenarios by 2050–2100 (NGFS Phase V, Nov 2024)

FOREWORD

What the Evidence Shows

The most useful thing this paper does is move the climate discussion out of the future tense.

Munich Re's annual report documented $320 billion in global economic losses from natural disasters in 2024, the fifth costliest year since 1980. Insured losses hit $140 billion, the third highest on record. But more important than the headline number is the trajectory behind it. Over the past decade, insured losses from natural disasters have grown at 5 to 7 percent annually in real terms, a compounding rate that is likely to double the financial burden over the coming decade.

The composition of those losses is changing as well. A growing share now comes from secondary perils: wildfires burning in places and seasons they historically did not, flash floods hitting cities whose drainage infrastructure was designed for different precipitation patterns, and severe convective storms producing hail damage at twice the historical insured loss rate. These are conditions where historical loss distributions are becoming less reliable as predictors of future exposure.

The operational consequence is that the gap between an institution's model-based risk estimates and its actual exposure is widening. Financial institutions are processing events shaped by environmental conditions that no longer behave consistently with the historical baselines many risk systems were trained on.

Barry J. Pruss
Executive Advisor, Enterprise AI & Operational Intelligence
Fulcrum Digital

Book a Call

CLIMATE RESILIENCE HAS BECOME AN ENGINEERING CHALLENGE WITH A DEADLINE. INSTITUTIONS THAT BUILD THE SENSING, MODELLING, AND RESPONSE INFRASTRUCTURE EARLY WILL HOLD AN ADVANTAGE OTHERS CANNOT ASSEMBLE UNDER PRESSURE.

This paper is written for financial services leaders now being asked to turn climate risk from a reported exposure into an operational capability. For AI leaders, that means building systems that can interpret live signals. For technology leaders, it means providing the infrastructure those systems depend on. For risk and architecture leaders, it means making sure the institution can act on those signals with governance, speed, and accountability. Every major claim in this paper is sourced because the conversation has moved beyond broad warnings. The work now sits in the systems themselves.

CHAPTER 1

The Operational Cost of Climate Disruption

Five recent climate and disaster events that exposed the operational limits of financial institutions under stress.

Financial institutions spent years discussing climate risk as a long-range exposure. The operational consequences arrived much faster. Between 2024 and 2025, floods, wildfires, earthquakes, and severe convective storms disrupted banking infrastructure, insurance operations, payment systems, and regional service continuity across multiple continents.

Valencia, Spain - October 2024: When Infrastructure Lagged Behind the Forecast

On 29 October 2024, a DANA weather system—a cold air mass isolated at altitude over warm Mediterranean waters—delivered more than 500 millimeters of rainfall to parts of Valencia province within 48 hours. Climate attribution research later found that the event was made roughly twice as likely by the 1.3°C of warming already experienced.

The flooding killed more than 237 people. For financial services, the operational consequences were immediate and measurable. Spain's central bank estimated that lenders carried approximately €20 billion in loan exposure to the affected region, equivalent to 1.8% of total loan volume across Spanish banks. Within days, 9% of ATMs in the flood zone were out of service and 37 of 298 bank branches remained closed. The Bank of Spain activated emergency payment system continuity measures and deployed real-time credit monitoring as infrastructure disruptions spread across the region. The Spanish government later released €14.37 billion in emergency relief, including ICO guarantee facilities and loan moratoria, after banking operations and payment access had already been affected.

KEY FINDING

Spain is considered one of Europe's most physically exposed banking markets for climate risk. The ECB's climate stress testing found that more than 60% of bank loans in Spain, Greece, and Portugal face elevated physical exposure, defined as more than a 1% annual probability of wildfire or flood impact. The Valencia flooding fell well within those modeled risk boundaries. (Sources: S&P Global Market Intelligence / ECB Climate Stress Test data, 2024)

The meteorological warning itself was not the failure point. The national weather agency had already issued red alerts before the flooding intensified. The breakdown appeared further downstream: in the coordination layer connecting forecast visibility, institutional escalation, operational response, and public communication. Japan's embassy in Spain, drawing on the same meteorological data, warned its citizens roughly a day before regional mobile alerts were issued locally.

For financial institutions, the lesson sits less in the forecasting systems themselves than in the operational infrastructure surrounding them. Environmental signals arrived on time. But the response systems handling escalation, communication, continuity coordination and regional operational visibility struggled to move at the same speed.

Japan - New Year’s Day 2024: The Compounding Failure

The Noto Peninsula earthquake struck at 16:10 local time on 1 January 2024, a public holiday when financial institutions were operating with reduced staffing and lower monitoring coverage. The magnitude-7.5 event triggered tsunamis, landslides, fires, and infrastructure failures across three prefectures. More than 36,000 households in Ishikawa and Toyama prefectures lost power. Financial and payment services were disrupted across multiple regions. Japan’s Financial Services Agency coordinated emergency countermeasures, including financial support lines and regional liquidity coordination for banks operating in the affected areas. The physical damage across Ishikawa, Toyama, and Niigata prefectures between ¥1.1 and ¥2.6 trillion ($7.5 to $17.6 billion USD).

A post-disaster analysis commissioned by the Japanese government identified a recurring operational limitation \during the response: institutions and emergency teams lacked a consolidated, real-time view of where affected individuals were located, which services remained available, and how conditions were changing across impacted areas. Ishikawa Prefecture later assembled what became described as Japan’s first disaster victim database using LINE messaging data and Suica IC card information, though the system took weeks to organize after the earthquake itself.

The operational strain for financial institutions in Noto came from the way multiple failures compounded simultaneously: reduced staffing, disrupted communications, regional service outages, and fragmented visibility into customer conditions across affected areas.

By the time emergency coordination systems were assembled, institutions were already operating inside the disruption rather than ahead of it.

Canada - 2024: The Year the Insurance Model Broke

The Insurance Bureau of Canada reported CAD $8.5 billion in insured losses from severe weather events in 2024—the highest annual total in Canadian history, nearly triple the 2023 figure and roughly 12 times the annual average recorded between 2001 and 2010 (CAD $701 million). Four catastrophic events during July and August alone generated over CAD $7 billion in insured losses and more than a quarter of a million insurance claims, approximately 50% more than Canadian insurers typically process across an entire year.

A hailstorm that struck Calgary on 5 August produced roughly CAD $3 billion in insured losses in little more than an hour. Flash flooding across the Greater Toronto Area caused nearly CAD $1 billion in insured damage. Remnants of Hurricane Debby triggered severe flooding in Quebec, generating an estimated CAD $2.7 billion in insured losses.

TD Economics later reported that Canadian P&C insurers incurred underwriting losses in personal property across both 2023 and 2024, paying out $1.01 for every $1 collected. In Alberta, the ratio was 1.20. Reinsurance premiums also climbed sharply, increasing 25 to 30% during 2023 renewals and 50 to 70% for portfolios with recent losses.

The Structural Protection Gap

Beyond insured losses, approximately $24 billion in additional uninsurable damage from Canada’s 2024 events was absorbed by governments, businesses, and individuals. Swiss Re estimates the global protection gap at $181 billion for 2024. The World Economic Forum places it at $223 billion. The gap compounds annually at roughly the same rate as insured losses, meaning the industry cannot model its way out of this through better pricing alone. The IBC publicly warned that parts of Canada could become uninsurable within a decade if action is not taken.

Australia - The Regulatory Evidence

The Australian Prudential Regulation Authority ran one of the most comprehensive bank-level climate vulnerability assessments undertaken by any major regulator in recent years. The 2022 Banking Climate Vulnerability Assessment engaged Australia's five largest banks under two warming scenarios: 1.8°C and 3°C+. Two of the five banks modelled lending losses reaching roughly three times the long-run historical average by 2050 under the higher scenario.

APRA followed with an Insurance Climate Vulnerability Assessment in 2023. The 2022 Australian floods had already provided a live stress test: more than 300,000 insurance claims totaling $7.4 billion in insured losses. APRA's Director of General Insurance and Banking later remarked that "billion-dollar insurance events should be expected almost yearly."

The Australian assessments are significant because they move beyond disclosure frameworks into quantitative operational modelling. Banks and insurers were required to examine how sustained warming conditions would affect lending exposure, insurance affordability, capital resilience, and long-term financial stability under increasingly severe physical risk scenarios.

The Global Pattern

Five jurisdictions. Different geographies, different disaster types, different regulatory systems. The operational pattern, however, is remarkably consistent: forecasting capability continues to improve, while institutional response systems struggle to operate at the same speed as the events themselves. This is the core systems problem that defines the climate resilience challenge for financial services.

Climate risk is now an operational condition financial institutions are expected to function through in real time.

Lessons Learned

- Forecast accuracy is improving faster than institutional response coordination.
- Catastrophe exposure now accumulates across infrastructure, staffing, payments, claims operations, and customer communication simultaneously.
- Event frequency is compressing operational recovery windows that insurers and banks historically treated as isolated incidents.
- Climate stress is beginning to surface as an operational continuity issue before it appears as a balance-sheet issue.
- Historical catastrophe assumptions are becoming less reliable as baseline environmental conditions shift.
- Regulatory attention is moving closer to operational resilience, continuous monitoring, and response capability under live conditions.
- Institutions operating with fragmented exposure visibility lose critical response time during rapidly evolving events.

CHAPTER 2

The Architecture of Predictive Intelligence Infrastructure

The infrastructure financial institutions are building, the capabilities they are missing, and why the gap persists.

Predictive intelligence begins before the model. A financial institution first has to capture live environmental signals, connect them to its exposure data, interpret what the event could mean, and carry that assessment into the workflows that govern response. Most investment still concentrates in the middle of that chain. The larger weakness sits at the edges: the quality of incoming signals and the institution's ability to act on what the system detects.

Layer One: Environmental Signal Ingestion

Signal ingestion is the data infrastructure that continuously pulls information from outside the institution's own systems. For climate operational intelligence, this spans weather and atmospheric data, satellite and aerial imagery, seismic sensor networks, hydrological monitoring, wildfire progression data from agencies like CAL FIRE or Canada's interagency fire monitoring systems, and IoT telemetry from built infrastructure.

Most large financial institutions already operate sophisticated signal ingestion infrastructure for financial market data, processing millions of events per second with subsecond latency. Extending that engineering discipline to environmental data is technically achievable. The larger constraint has been investment prioritization. Institutions capable of processing real-time FX data at microsecond precision may still update physical-risk exposure across mortgage portfolios on a quarterly basis.

The commercial catastrophe modeling industry has built the most advanced version of signal ingestion for insurance applications. Moody's RMS now operates high-definition models at up to one-meter grid resolution for flood and wildfire risk, incorporating vegetation moisture content, soil composition, and drainage capacity. These models are steadily moving from periodic batch analysis toward near-continuous event monitoring.

PRODUCTION EVIDENCE

After the January 2025 Los Angeles wildfires, Moody's AI-powered image analysis compared pre-event satellite imagery with post-event data to classify every affected structure as fully destroyed, partially damaged, or untouched—distinguishing primary residences from outbuildings at the parcel level. The analysis was completed within hours of the fire perimeter stabilizing, compared to days or weeks for traditional ground-based assessment. The same methodology had been applied after Hurricane Ian's Florida landfall in September 2022. (Source: Moody's RMS, February 2026)

Layer Two: Predictive Reasoning

The traditional catastrophe (CAT) model was built on stationarity: the assumption that the future distribution of events would resemble the past. That assumption is becoming less reliable. Atlantic hurricane intensification rates have increased 25% over four decades. Wildfire seasons that historically lasted four months now extend closer to eight. Flood events classified as 100-year occurrences have struck the same regions multiple times within a decade. The Valencia DANA event was made roughly twice as likely by current warming levels.

THE STATIONARITY PROBLEM

Traditional CAT models depend on historical loss distributions. As those distributions shift for wildfire, flood, and severe convective storm perils, the gap between modeled risk and realized loss widens. Verisk became the first company to submit a wildfire catastrophe model for ratemaking approval in California after new regulation took effect on January 2, 2025. California had previously only permitted CAT models for earthquake risk. The industry's long-standing dependence on historical wildfire patterns had become untenable. (Source: CAS Actuarial Review, 2025)

Machine learning models address part of this challenge by incorporating real-time inputs alongside historical training data. For wildfire, ML-enhanced models now integrate satellite-derived vegetation moisture content, live wind conditions, power-grid infrastructure proximity, and ignition clustering patterns. Research has shown that these models can forecast wildfire spread several hours ahead of the fire front with measurable improvements over purely physics-based approaches in complex terrain.

Flood modeling has evolved in a similar direction. Real-time river gauge data, soil saturation measurements from IoT sensors, and high-resolution precipitation nowcasting support probabilistic flood trajectory forecasting at city-block level.

The governance dimension of AI-powered predictive modeling remains a major operational gap for many financial institutions. A model trained primarily on historical climate patterns is already operating on delayed assumptions. Models continuously calibrated against current sensor data operate much closer to the conditions they are attempting to predict. Maintaining that capability requires an MLOps discipline—model versioning, drift detection, retraining pipelines, and performance monitoring—that many financial institutions did not prioritize before AI systems became operational dependencies.

Layer Three: Operational Response Integration

The third layer is where many institutions have the largest capability gap and where the operational consequences become most visible. This is the layer that translates predictive output into operational action: client communications, credit exposure adjustments, continuity procedures, payment rerouting, claims mobilization, or escalation to human judgment when conditions exceed automated thresholds.

Parametric insurance provides a working example of how the response integration layer can operate. Unlike traditional indemnity insurance, parametric contracts trigger payouts automatically when a pre-defined physical threshold is crossed—wind speed, rainfall accumulation, seismic intensity, or river gauge height. The payout can arrive within days because the assessment logic is already embedded into the contract structure.

The parametric market is projected to reach $34.4 billion by 2033 (WEF). Munich Re is developing a parametric flood insurance pilot with the Mississippi River Cities and Towns Initiative. The UNDP has documented the Reserve Bank of Fiji's national parametric micro-insurance program—the first central bank globally to lead such an initiative—with more than 1,700 beneficiaries receiving payouts within days of qualifying rainfall events.

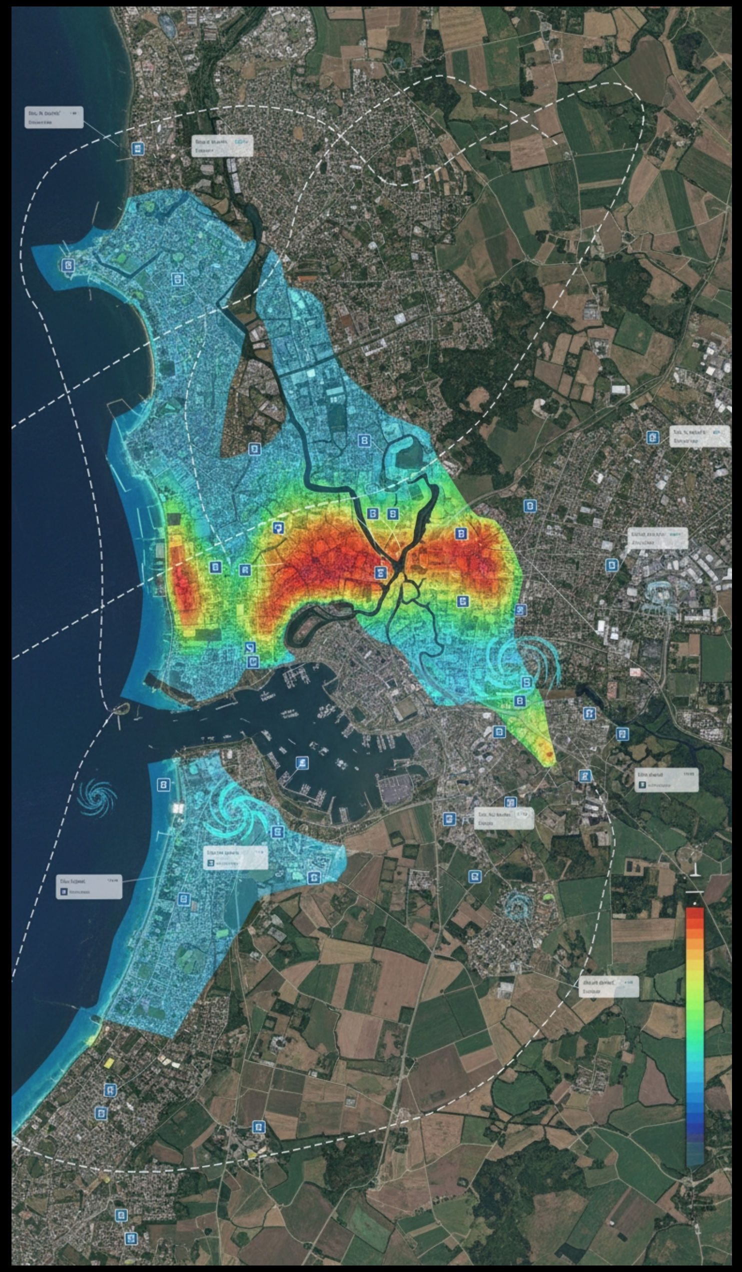

Where the Investment Gap Sits

LAYER	FUNCTION	CORE CHALLENGE	CURRENT STATE
Signal Ingestion	Real-time environmental data	Data engineering, pipeline governance, latency	Strong for market data; weak for physical environment data; gap is deliberate investment, not technical limitation
Predictive Reasoning	AI/ML forecasts and exposure estimates	Model stationarity, training data recency, continuous calibration	Significant investment; most models still run periodically; stationarity assumption increasingly unreliable
Response Integration	Forecasts translated to operational action	Authority structures, human-AI governance, escalation design	The weakest layer; forecast exists but is not wired into operational decision-making with appropriate speed and accountability

Lessons Learned

- → Predictive intelligence depends as much on data pipelines and response workflows as it does on model quality.
- → Environmental data needs to be governed with the same seriousness financial institutions already apply to market data.
- → Physical-risk exposure loses value when it is updated quarterly while the underlying event is moving hourly.
- → Historical loss data is becoming a weaker baseline for wildfire, flood, and severe storm modeling.
- → Continuous model monitoring is now part of resilience infrastructure, not a technical afterthought.
- → The response layer is usually where predictive capability loses operational value.
- → Predefined triggers, audit trails, and escalation rules determine whether predictive intelligence can be trusted during live events.

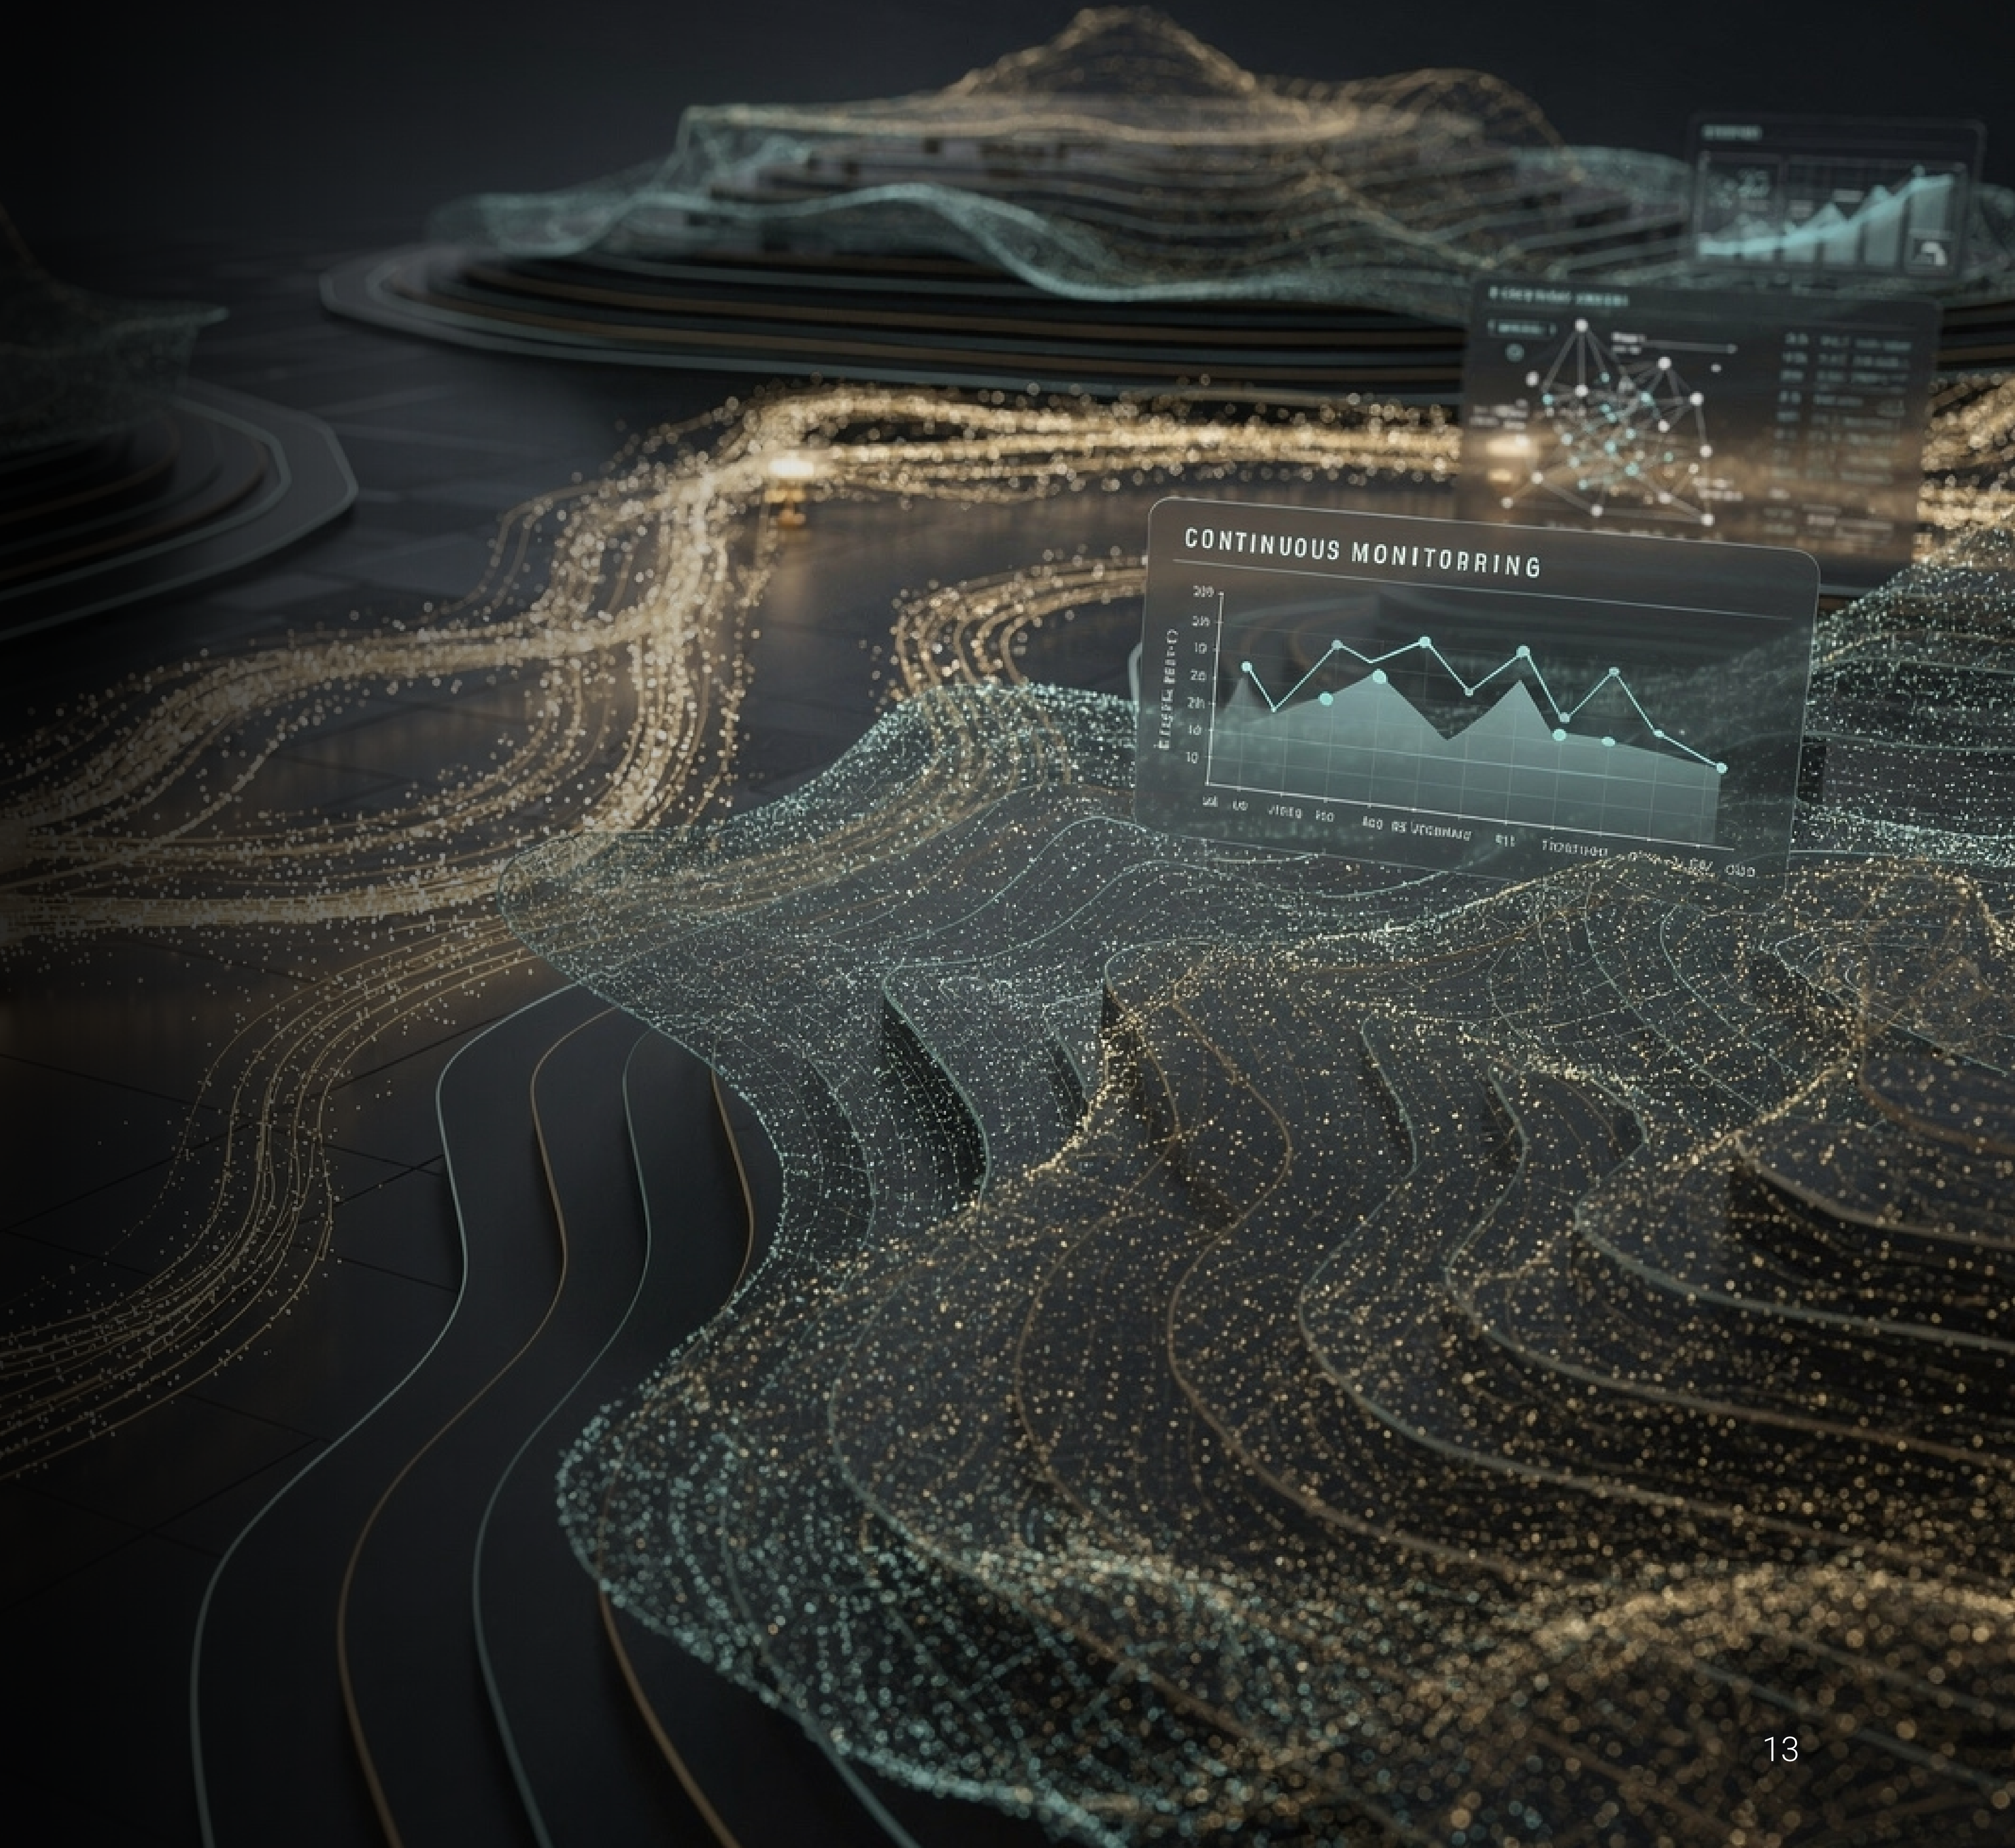

CHAPTER 3

Quantum Computing and the Limits of Readiness

An honest account of where quantum methods contribute
to financial services resilience and where they remain experimental

Quantum computing occupies an unusual position in the financial services technology landscape. The commercial quantum future—general-purpose fault-tolerant quantum computers delivering transformative advantages across broad problem classes—is genuine as a long-horizon capability. As a near-term operational resource, it requires careful delineation. This chapter is an honest attempt at that delineation.

What the Research Shows

The academic and actuarial literature on quantum applications in financial services has expanded since 2023. A 2024 paper published on arXiv (Liu et al.,arXiv:2410.20841) constructed a systematic review of connections between quantum computation and insurance problems, including experimental demonstrations of variational quantum algorithms for reinsurance type allocation problems.
A March 2026 piece in The Actuary Weekly examined quantum computing (QC) in catastrophe modeling for life and health insurance, noting that the practical value for actuaries "lies in its ability to model high-dimensional probability distributions more efficiently than current methods allow", particularly for cascading risks across multiple timescales. The paper explicitly recommends incremental proof-of-concept adoption while cautioning that fault-tolerant QC remains "horizon technology."

HONEST ASSESSMENT

The consensus across current actuarial and computational research: (1) quantum annealing and variational algorithms are demonstrating meaningful advantages on specific optimization problems today, including reinsurance portfolio allocation and constrained resource dispatch; (2) quantum simulation methods are relevant to physical systems modelling including seismic stress accumulation; but (3) fault-tolerant general-purpose quantum computing at the scale required for most financial services applications remains 3–10 years from commercial availability. Any institution claiming broad production quantum capability today should be evaluated carefully. (Sources: arXiv 2410.20841; MDPI Risks 2023; The Actuary Weekly, Mar 2026)

Where Quantum-Adjacent Methods Are Contributing Now

For financial services resilience, the highest near-term value from quantum or quantum-inspired methods is in simulation rather than general computation. Quantum-inspired algorithms, which borrow quantum algorithmic structures but run on classical hardware, can improve solution quality for specific classes of high-dimensional simulation problems within operationally feasible compute times. For reinsurance optimization and portfolio allocation under catastrophic uncertainty, variational quantum algorithms have demonstrated competitive performance against classical optimization methods on benchmark problems. The commercial relevance is clearest for large-scale multi-peril, multi-geography portfolio management where the optimization space is combinatorially large.

The Seismic Intelligence Case

The intersection of quantum methods and earthquake prediction illustrates both genuine scientific interest and current operational limitations. Seismic stress accumulation models require high-resolution simulation of interactions between tectonic stress fields, fault geometry, groundwater conditions, and surface topography. Classical simulation approximates these interactions by discretizing the system. Quantum algorithms designed for complex physical systems simulation offer a theoretically grounded pathway to improving temporal and spatial accuracy of seismic forecasting. The research direction is already visible in national geological work and in financial institutions with material seismic exposure.

The Infrastructure Dependency

Quantum enhancement of operational intelligence systems requires the classical AI infrastructure to be built first. Quantum methods provide computational advantage on specific algorithmic sub-problems within a system that primarily runs on classical infrastructure. Early advantage will sit with institutions already building production-quality classical predictive infrastructure.

QUANTUM CAPABILITIES WILL ADD THE MOST VALUE WHERE THE DATA FOUNDATION, AI INFERENCE LAYER, AND RESPONSE ARCHITECTURE ARE ALREADY MATURE. WITHOUT THOSE FOUNDATIONS, QUANTUM REMAINS AN ADVANCED COMPUTATION LAYER WITH NOWHERE MEANINGFUL TO OPERATE.

Lessons Learned

- Quantum computing remains a long-horizon capability for most financial services applications.
- Near-term value is more likely to emerge through targeted modeling and simulation problems than broad operational replacement.
- Climate forecasting, seismic modeling, and catastrophe simulation are attracting serious quantum research because of their computational complexity.
- Classical AI infrastructure still determines whether predictive systems can operate reliably under live conditions.
- Institutions without mature data pipelines and response architecture are unlikely to benefit meaningfully from quantum enhancement.
- The strongest near-term advantage may come from hybrid systems that combine classical AI with selective quantum methods.
- Financial institutions evaluating quantum systems need clearer distinctions between active research, pilot capability, and production readiness.

CHAPTER 4

The Infrastructure of Resilience

The capabilities, coordination systems, and infrastructure patterns appearing inside resilient financial institutions

Resilience now depends on how well financial institutions connect sensing, analysis, decision authority, and response capacity before disruption begins. Climate events have already shown the cost of delay, and quantum-enhanced methods may improve parts of the forecasting layer over time. The immediate requirement is more practical: institutions need systems that can detect changing exposure, propagate scenarios, activate response workflows, and escalate decisions while conditions are still moving.

The Four Capabilities That Make the Difference

1. Continuous Exposure Monitoring

The ability to see operational exposure under current physical conditions, rather than the assumptions from the last model update. The Bank of Spain activated real-time credit monitoring after Valencia; this architecture brings that visibility earlier in the response cycle.

2. Near-Real-Time Scenario Propagation

The ability to translate a developing event into updated exposure across the loan book, investment portfolio, and operational footprint over the next 24, 48, and 72 hours, using live data feeds and scenario updates measured in minutes.

3. Pre-Configured Operational Playbooks

Response workflows triggered when system-generated signals cross a material exposure threshold. Client communication, branch continuity, liquidity coordination, or credit adjustment protocols begin without waiting for manual assessment. Valencia and Noto showed the cost of assembling response after disruption is underway.

4. Human Escalation Architecture

The governance structure that determines when automated response is sufficient and when human judgment is required. The ECB's 2025 stress testing direction reinforces the need for auditability: recommendations must be traceable to the underlying data and decision path.

The Cloud Infrastructure Underneath This

The operational intelligence architecture described above places specific demands on the infrastructure supporting it: variable compute intensity, latency sensitivity, and governance complexity. Cloud migrations completed by several major financial institutions over the past three years are relevant here less for cost reduction than for the computational ceiling they removed. A risk computation platform previously constrained to 5,000 cores—with incomplete analysis runs becoming a real operational risk during periods of high demand—can after cloud migration scale to 80,000 cores on demand, enabling continuous scenario analysis that earlier infrastructure constraints made impractical.

The migration itself does not solve the resilience challenge. It expands the operational range within which the challenge can be addressed.

RESILIENCE MATURITY	DESCRIPTION	RESPONSE TIMELINE	CURRENT STATUS
Reactive Continuity	Post-event response following pre-written plan	Hours to days	Most institutions; response begins after operational damage is occurring
Model-Informed Risk Management	Pre-event capital/ pricing adjustment from modelled probabilities	Days to weeks	Standard Tier-1 practice; does not adapt to developing events in real time
Operational Intelligence	Continuous signal monitoring with automated response triggering	Minutes to hours	Leading institutional practice; requires all three connected infrastructure layers
Adaptive Resilience	Systems incorporating real-time event data into operational decisions continuously	Near-real-time	Emerging capability at select institutions; quantum-enhanced simulation contributes in specific high-complexity scenarios

The Australian Banking Industry Model

Australia's response to climate-driven continuity requirements offers a useful operational example. The Australian Banking Association maintains standing hardship response protocols for disaster events, with dedicated teams activated at the point of a declared natural disaster rather than assembled in response to one. When Tropical Cyclone Alfred made landfall in early 2025, APRA's Director acknowledged that 75,000 insurance claims had already been lodged before the event had fully resolved. The 2022 floods, which generated 300,000 claims and $7.4 billion in insured losses, had already become what APRA described as a "wake-up call" for the scale of natural disaster impact the industry should expect.

The ABA's pre-committed hardship response model demonstrates how continuity infrastructure can be established before disruption begins rather than constructed during active response.

Australian Banking Association

Lessons Learned

→ Resilience depends on how quickly institutions can connect detection, decision-making, and response.

→ Exposure visibility loses value when operational systems cannot react at the same speed as the event.

→ Static continuity plans degrade quickly during fast-moving climate disruption.

→ Response workflows become more effective when escalation thresholds are predefined before the event begins.

→ Real-time scenario propagation is becoming a core operational capability for banks and insurers with physical-risk exposure.

→ Human escalation logic needs to be embedded directly into automated decision systems.

→ Infrastructure limits can become resilience limits when compute capacity cannot support continuous analysis under stress.

→ Institutions with preconfigured response teams recover coordination time during live disruption.

CHAPTER 5

What Six Global Regulators Have Concluded

OSFI, Bank of England, ECB, NGFS, APRA, and the Basel Committee—converging on the same assessment about operational architecture

Six major supervisory bodies, operating across different legal systems and oversight traditions, have arrived at substantially the same conclusion: climate-related physical risks require financial institutions to build operational infrastructure to manage them credibly.

OSFI Guideline B-15 (Canada) Now in Effect

OSFI's Guideline B-15 entered force for Canada's Domestic Systemically Important Banks and internationally active insurance groups at fiscal year-end 2024, following one of the most extensive consultation processes in OSFI's history with more than 4,300 submissions. All other federally regulated financial institutions follow at fiscal year-end 2025.

B-15's operational requirements are more substantive than most discussion of the guideline acknowledges. The guideline requires each institution to "incorporate the implications of climate change… in its business model and strategy"; to "continuously monitor and prudently manage climate-related risks to remain financially and operationally resilient"; and to develop processes to "measure the current and potential future impact of climate-related risks on its portfolio of exposures over short, medium and longer term." The word "continuously" is not incidental. It reflects OSFI's assessment that climate risk management based on periodic model updates is insufficient relative to how physical climate events develop.

Bank of England - Climate Biennial Exploratory Scenario

The Bank of England's 2021 Climate Biennial Exploratory Scenario was the first major regulatory exercise to subject financial institutions to quantitative climate stress testing across three scenarios: early, late, and no additional policy action. The exercise covered both bank and insurance sector participants across physical risk, transition risk, and their interactions. The BoE has continued developing its climate risk toolkit, with the November 2024 Financial Stability Report reaffirming commitment to embed climate risks in the stress-testing framework as part of ongoing financial stability work.

ECB - Systematic Capital Integration

The European Central Bank began applying additional capital requirements to banks that fail to effectively manage climate and environmental risks in November 2022 and set staggered deadlines for full integration into ICAAP and stress testing by end of 2024. The ECB's 2025 EU-wide stress test integrated climate risk scenarios alongside standard adverse economic scenarios for the first time. Results were quantified: acute physical risk is estimated to add approximately 77 basis points of CET1 depletion beyond the standard adverse scenario. The ECB also found that "banks most exposed to climate-related losses may differ from those identified as most vulnerable in the broader EU-wide assessment", meaning climate risk exposure does not map cleanly onto traditional stress test risk concentration patterns.

NGFS Phase V Scenarios - The Scale of the Revision

The Network for Greening the Financial System published Phase V of its long-term climate macro-financial scenarios in November 2024. The revised damage function now estimates economic losses equivalent to 15% of global GDP by 2050 under 2 degrees of warming and 30% of GDP by 2100 under 3 degrees, three times higher than earlier NGFS estimates. These are the reference scenarios used by the Bank of England, ECB, APRA, and most other major supervisors for climate stress testing.

APRA - Banking and Insurance CVA

APRA's 2022 Banking Climate Vulnerability Assessment and 2023 Insurance CVA together represent the most comprehensive multi-institution quantitative assessment of climate vulnerability undertaken by any prudential supervisor. Two of Australia's five major banks projected lending losses triple the historic average by 2050 under the higher scenario. APRA subsequently characterized the scale of natural disaster impact now expected as requiring the industry to plan for "billion-dollar events almost yearly"—a reframing that treats climate disruption as an operational planning assumption rather than a stress test tail scenario.

What Six Regulators Are Collectively Saying

Read together, these frameworks share a common assessment: climate risk management based on periodic modelling and annual disclosure is insufficient. B-15’s “continuously monitor” expectation. The BoE’s “embed in stress testing” direction. The ECB’s capital requirements for failure to manage climate risk. The NGFS’s 3x upward revision to macroeconomic damage estimates. The APRA CVA’s finding of triple-average lending losses at two major banks.

Each is a different expression of the same institutional judgement: that institutions managing climate risk operationally will remain financially and operationally resilient as event frequency continues to rise.

Governance and the AI Trust Architecture

Every regulatory framework in this chapter requires that institutions demonstrate how they manage climate risk: through auditable processes, explainable models, and accountable decision structures. This requirement applies with particular force to AI-assisted operational systems. OSFI’s contemporaneous Guideline E-23 on Model Risk Management (effective May 2027) requires enterprise-wide AI/ML model inventories, explainability and fairness documentation, and lifecycle governance. The governance architecture for climate-related operational AI is not a compliance add-on. It is the accountability structure that makes the system trustworthy enough to act on in a high-stakes situation.

Lessons Learned

- → Climate risk supervision is moving steadily from disclosure requirements toward operational capability requirements.
- → Regulators across multiple jurisdictions are converging on continuous monitoring as a baseline expectation.
- → Physical-risk exposure is beginning to influence stress testing, capital treatment, and supervisory assessment simultaneously.
- → Climate-related losses are no longer being modeled as isolated tail events in major regulatory frameworks.
- → Institutions with weak operational visibility may face growing regulatory pressure even when traditional financial indicators remain stable.
- → AI-assisted climate systems require governance structures that can explain, trace, and defend operational decisions under review.
- → Model inventories, auditability, and lifecycle governance are becoming foundational requirements for enterprise AI deployment in regulated environments.

CHAPTER 6

Building the Sensing Layer

23

For institutions at different maturity stages, the logical order of investment and the failure modes to avoid

The Most Common Failure Mode: Starting With the Model

A consistent failure pattern emerges across production AI deployments in financial services environments: beginning with the predictive model rather than the data infrastructure. Organizations invest in AI capability before the data that model needs is clean, governed, and accessible at the temporal resolution that operational use requires. For climate operational intelligence, this means building a wildfire spread prediction model before establishing real-time vegetation moisture data feeds. Or building a flood trajectory model before integrating real-time river gauge data. The resulting system produces outputs too stale to drive operational response; the model's technical capability substantially exceeds its operational value.

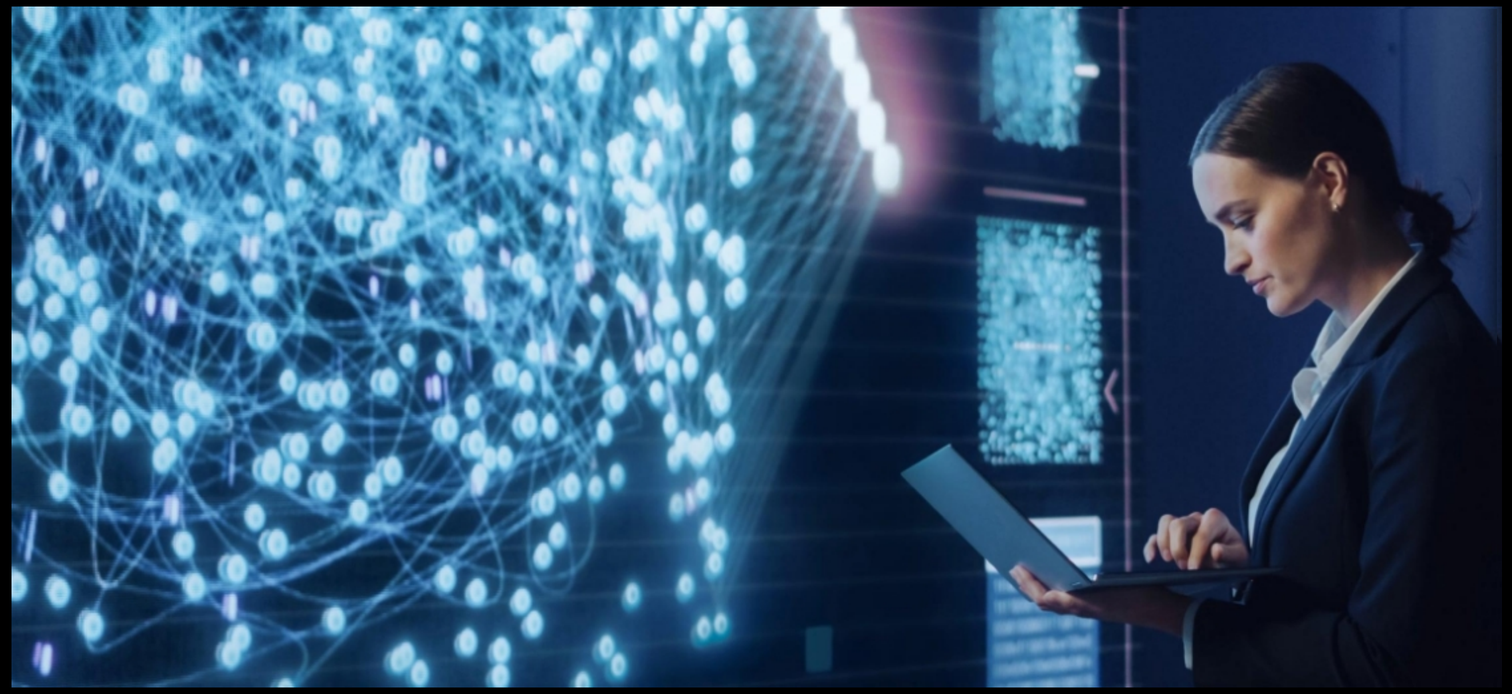

Start With the Data Foundation

The data engineering investment that makes operational intelligence possible is less visible than the AI model investment and requires sustained commitment to infrastructure that looks like overhead until the event that proves it was load-bearing. The foundation has three components: external signal integration (pipelines bringing environmental data into the institution's governed infrastructure with appropriate latency and quality controls); the portfolio data layer (a real-time view of current exposure across geographies and asset classes with sufficient geographic granularity); and the operational data layer (real-time status data on branch availability, ATM functionality, payment system capacity, staff deployment).

Build the Predictive Layer for Continuous Operation

Once the data foundation exists, the predictive layer should be built for continuous operation rather than periodic reporting. This requires an ML operations discipline that differs from traditional model development: models must be versioned, monitored, retrained when performance degrades, and rolled back when updates create unexpected behavior. Predictions must be logged with their inputs so the system's reasoning is auditable—a requirement that both operational trust and regulatory compliance demand. The catastrophe modeling firms provide the most sophisticated external reference point for what this looks like at scale: Moody's RMS operates over 400 models across nearly 100 countries, continuously updated with new data and research.

Wire It into Operations Last But Plan for It First

The operational response integration layer requires governance decisions that cannot be made by a technology team alone. Which operational responses can the system trigger automatically? What client communications can be initiated without human approval? At what exposure threshold does the system escalate to the Chief Risk Officer? These decisions require active engagement from CROs, COOs, and CCOs as participants in the design from the start. The governance architecture—explainability requirements, audit trail, and human review thresholds—is the accountability infrastructure that makes the system trustworthy enough to act on in a high-stakes situation. A system whose recommendations cannot be traced back to the data that generated them will not be trusted, and a system that is not trusted will not be used, regardless of its technical accuracy.

Where Fulcrum Digital Operates

Fulcrum Digital's position in this architecture is at the engineering layer between frontier research and production operational systems. Catastrophe modeling firms advance predictive science. Cloud infrastructure providers supply the compute platform. Research institutions advance quantum algorithms and physical simulation methods. The capability most consistently under-resourced—and most frequently determining whether the architecture becomes operational or remains experimental—is the production engineering connecting hose layers into systems that financial institutions can govern, maintain, and depend on.

Across more than 4,500 production engagements in financial services, insurance, and adjacent sectors since 1999, the failure patterns we encounter are concentrated in the production engineering layer: data pipelines built for development but not for operational resilience, AI operations infrastructure excluded from the original project scope, and governance architecture designed as a compliance exercise rather than as the accountability infrastructure the organization genuinely needed. The technical solutions are known. They have been built and are in production at financial institutions operating under regulatory constraints analogous to those described in this paper. The engineering discipline that keeps them working 12 months after deployment, in conditions different from the development environment, is the capability Fulcrum Digital brings.

Lessons Learned

- → Environmental intelligence depends first on data engineering, not model sophistication.
- → Predictive systems lose value when the data arrives too slowly for operational use.
- → Real-time exposure visibility requires external signals, portfolio data, and operational status data to work together.
- → Systems built for periodic reporting struggle during fast-moving events.
- → Model monitoring, rollback controls, and audit logging become critical once predictions influence operations.
- → Escalation, approval, and automated response rules need executive ownership from the start.
- → Production engineering often determines whether an AI pilot becomes a reliable operating system.

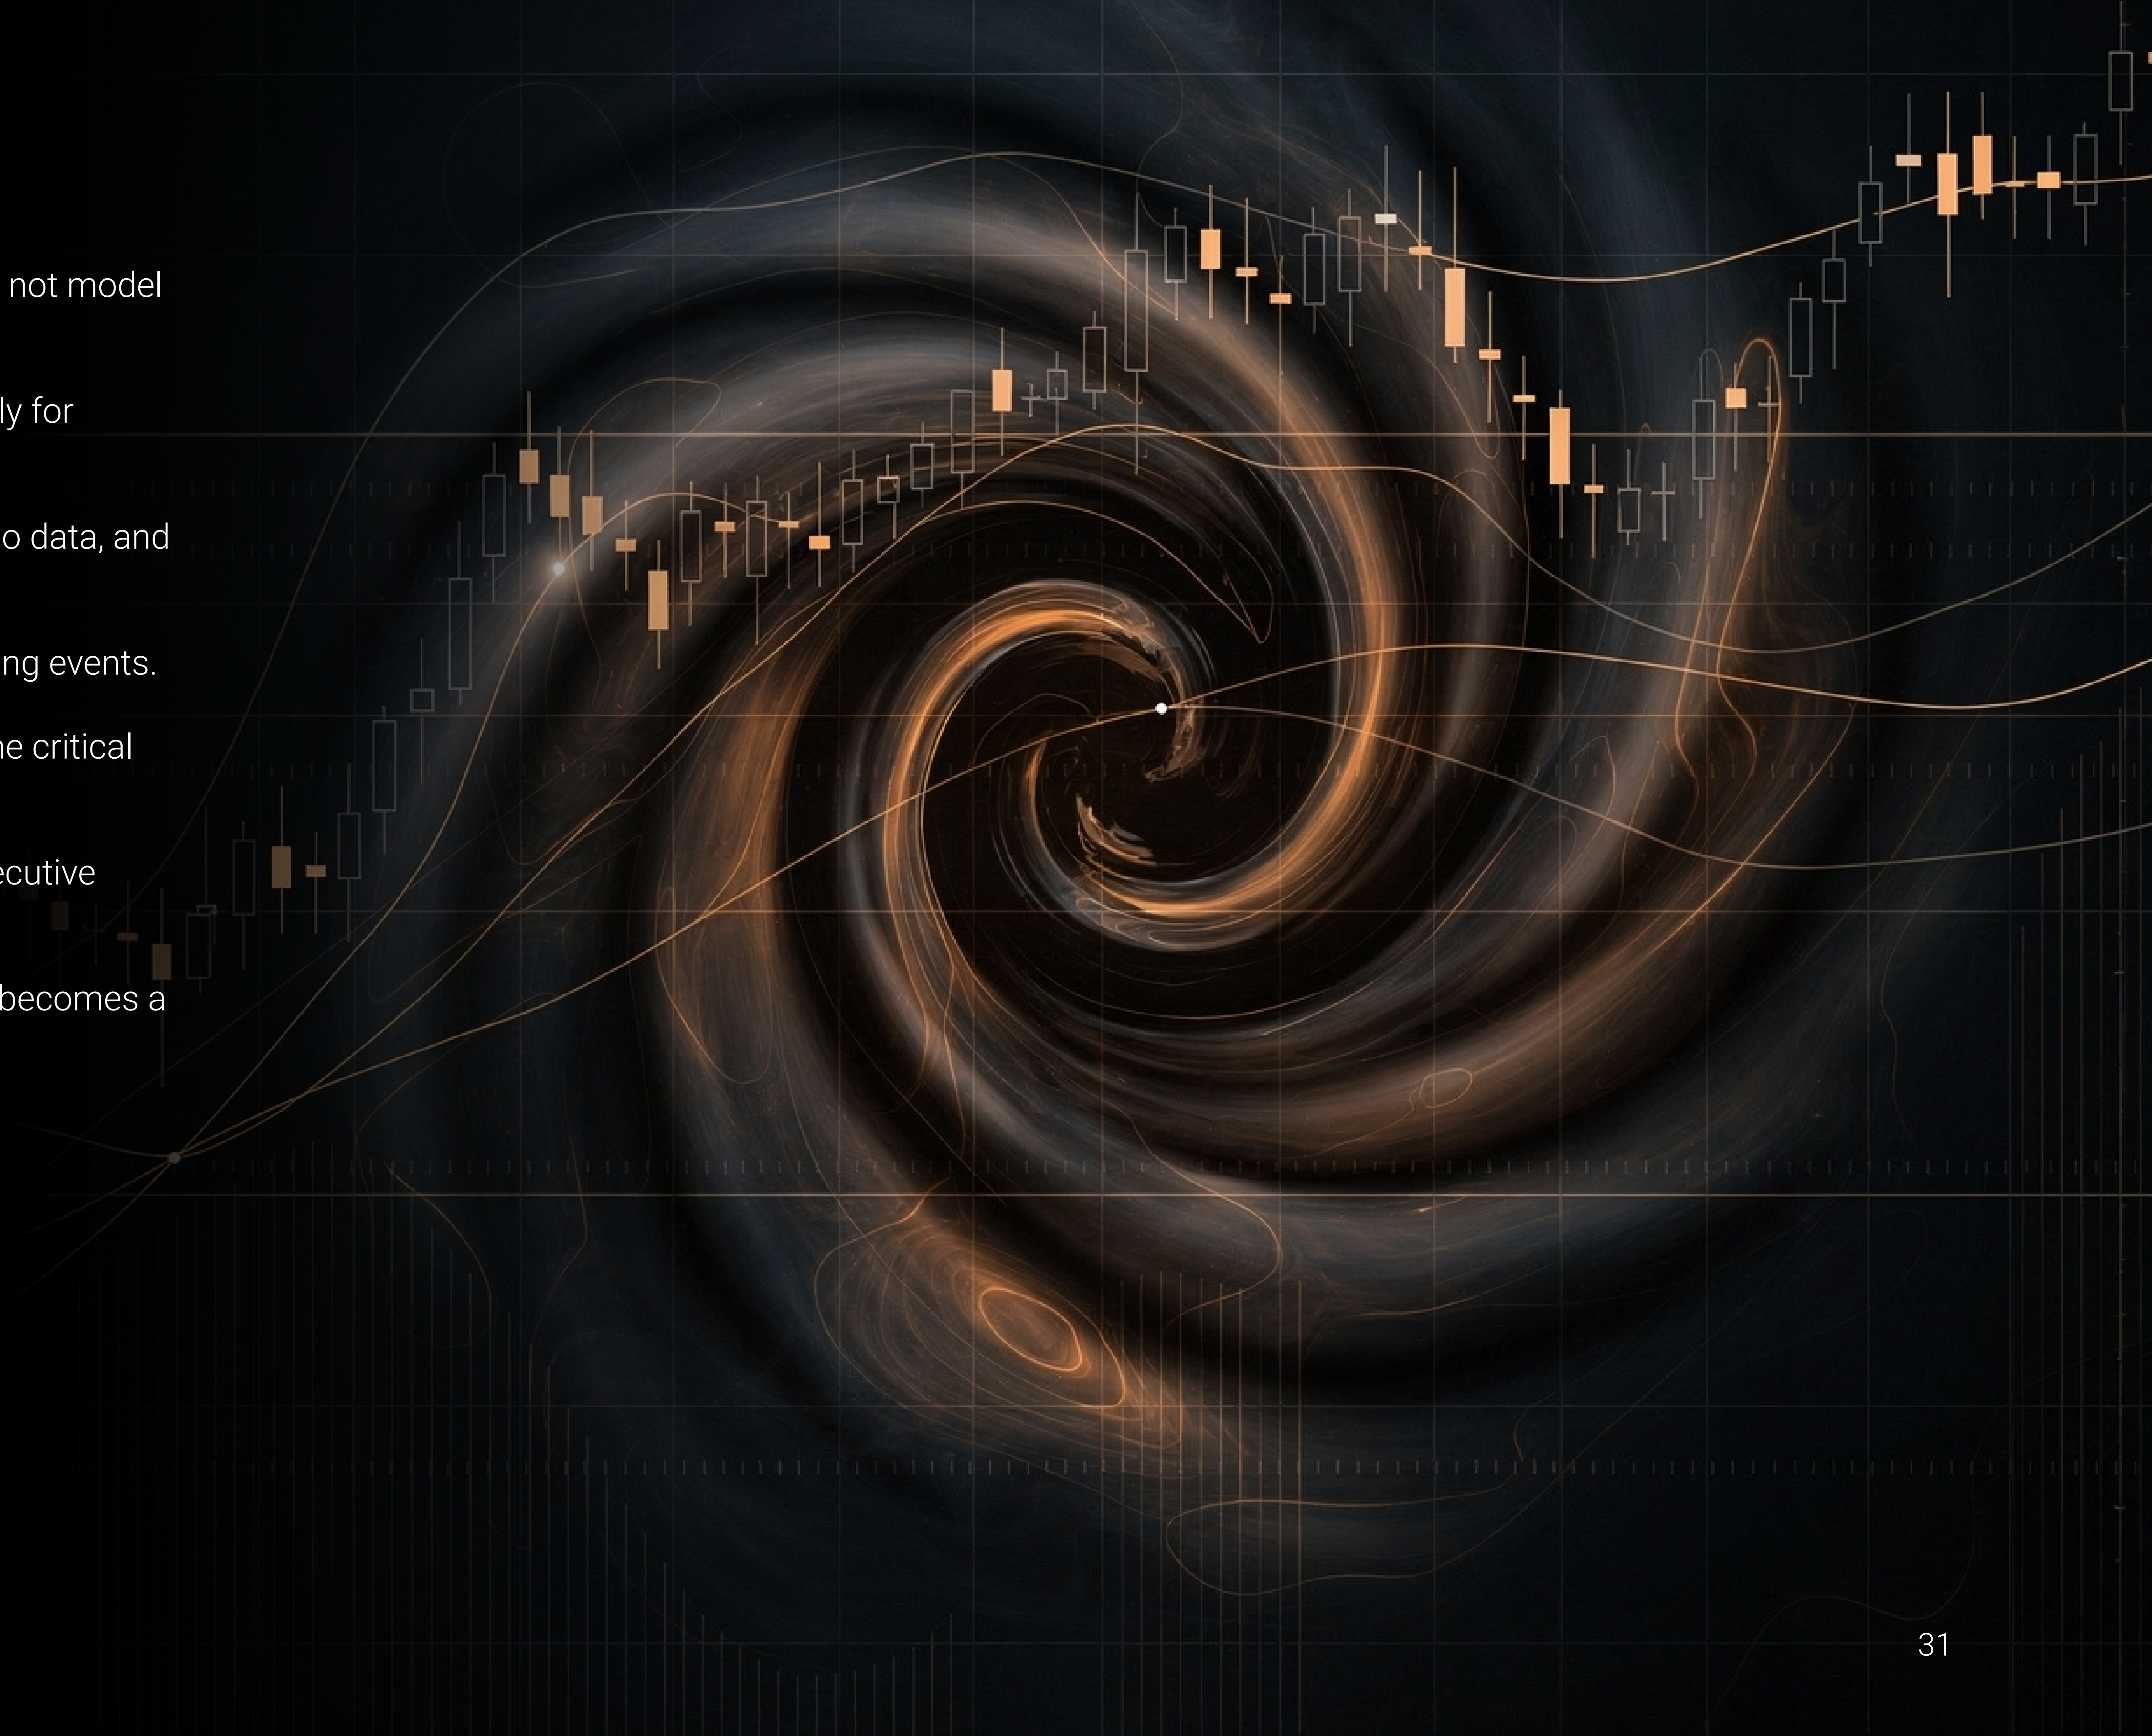

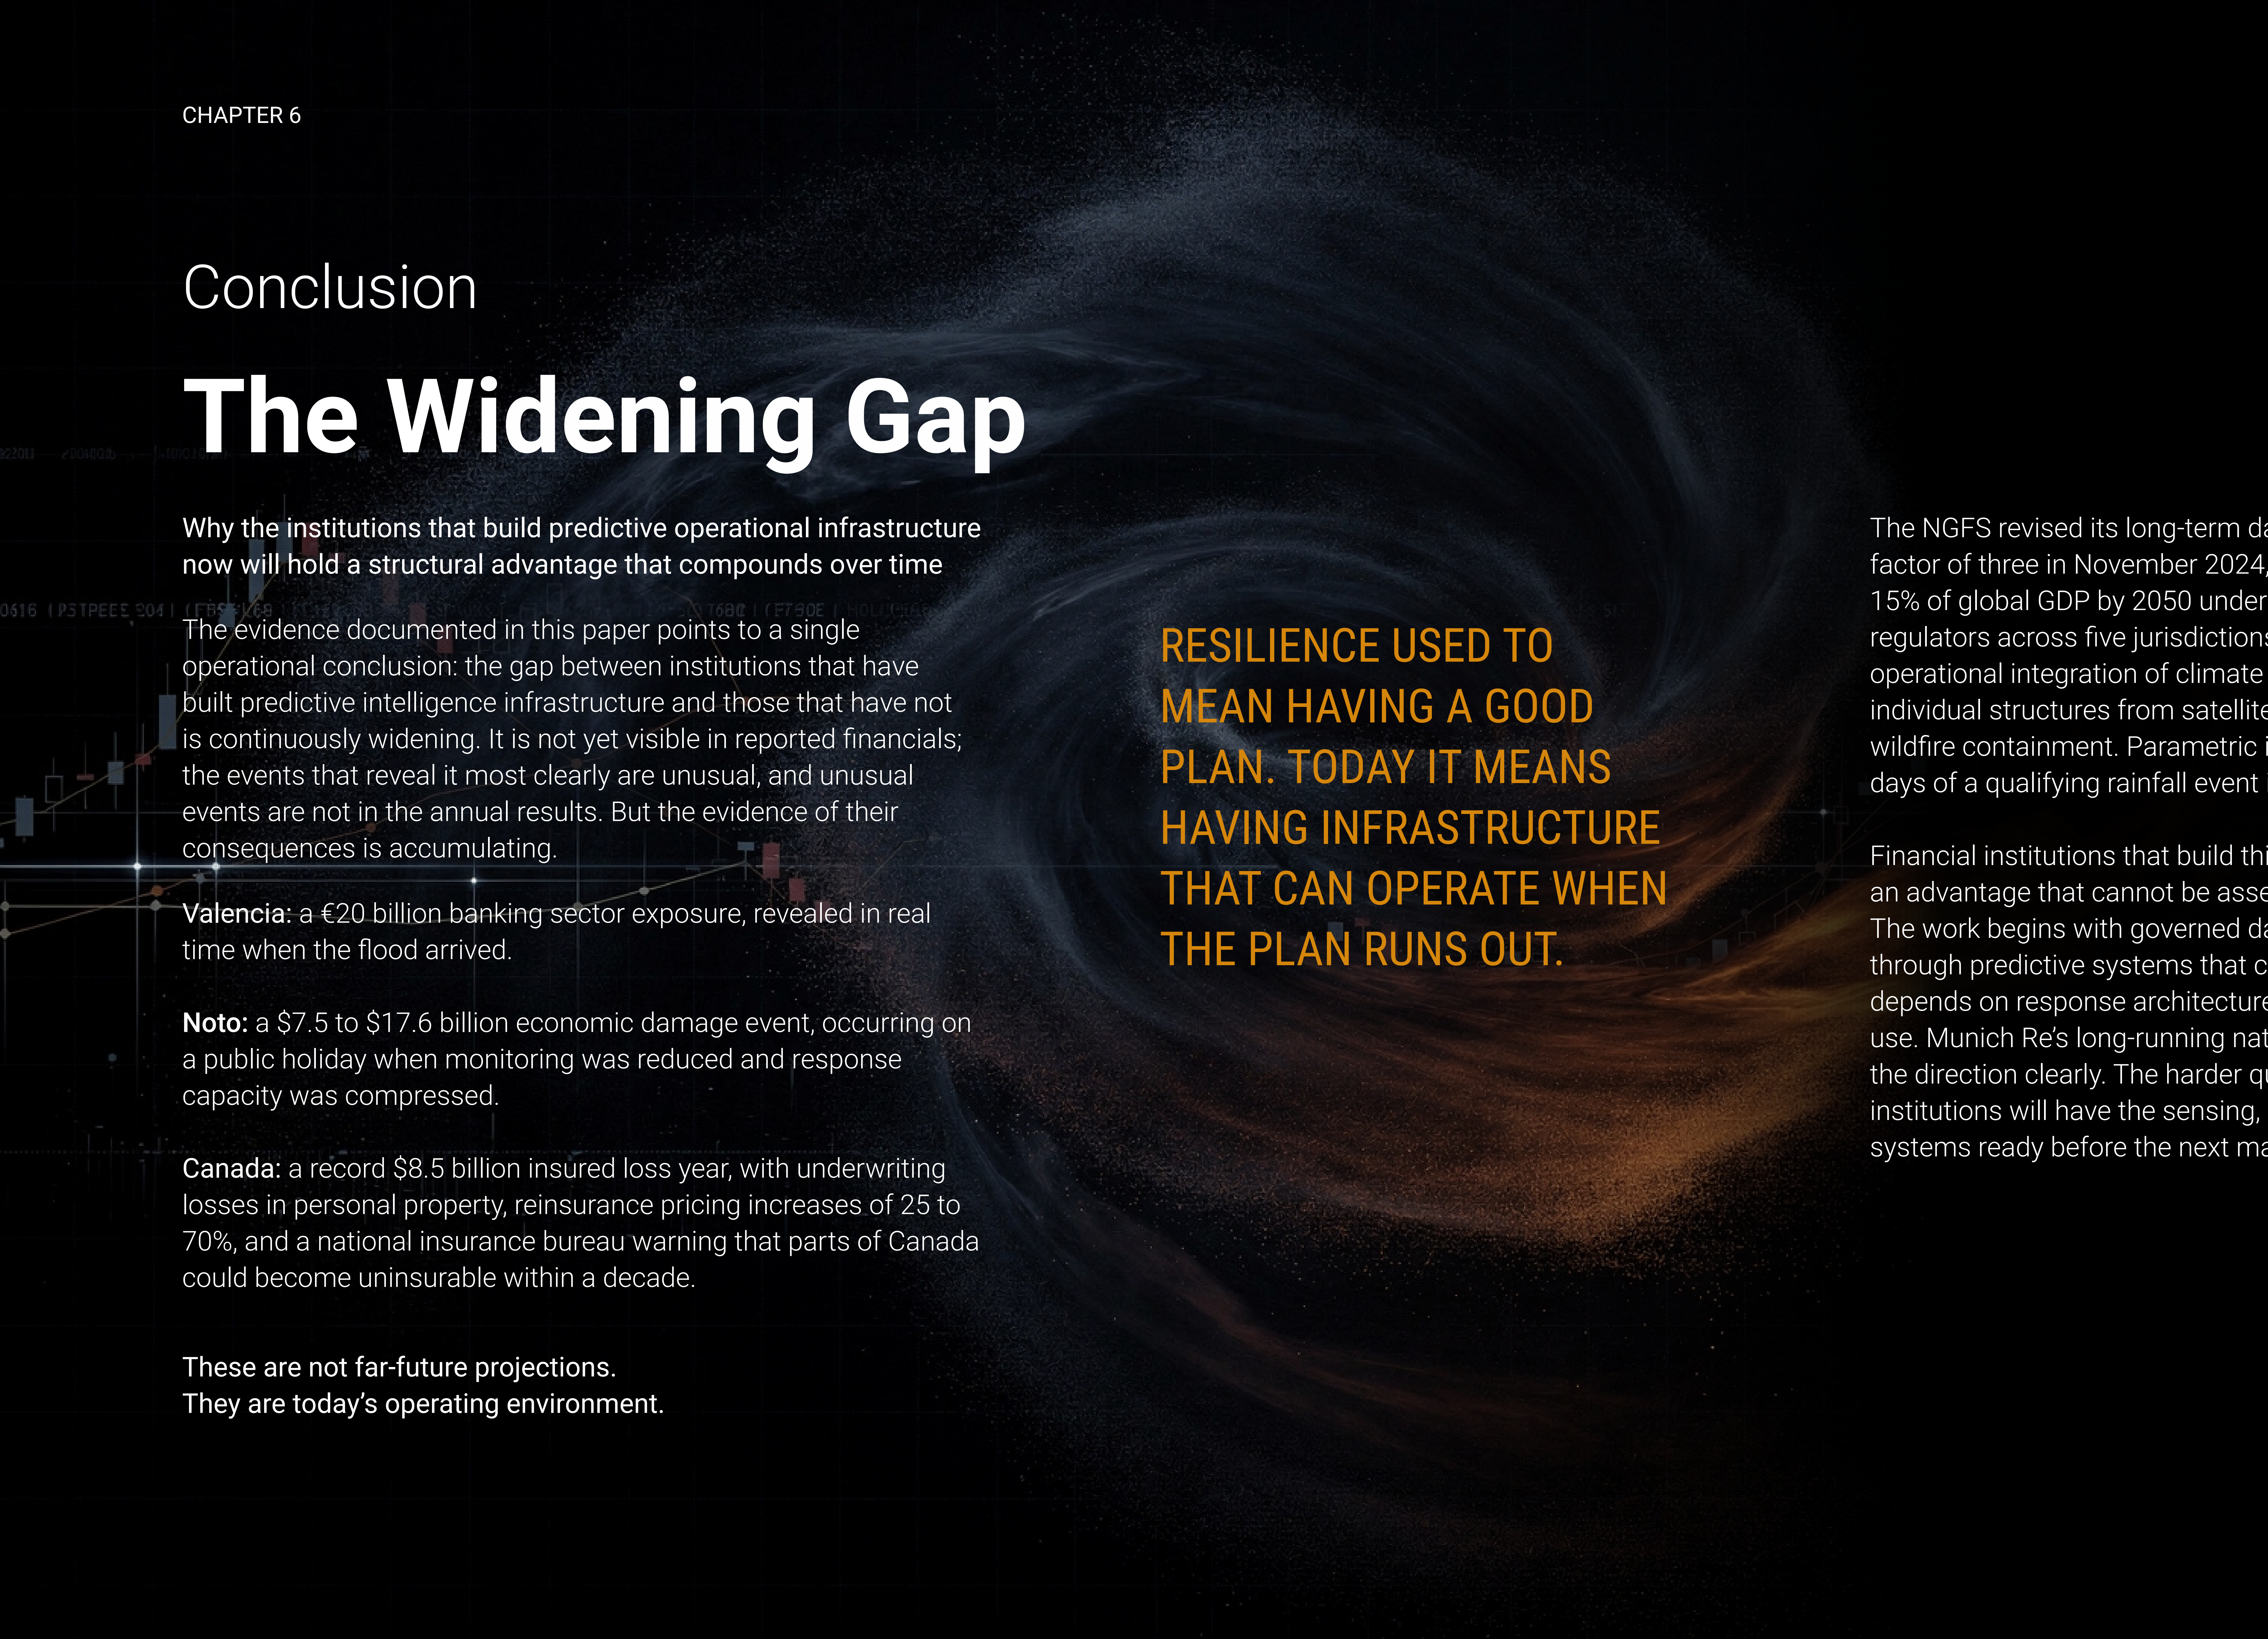

Conclusion

The Widening Gap

Why the institutions that build predictive operational infrastructure now will hold a structural advantage that compounds over time

The evidence documented in this paper points to a single operational conclusion: the gap between institutions that have built predictive intelligence infrastructure and those that have not is continuously widening. It is not yet visible in reported financials; the events that reveal it most clearly are unusual, and unusual events are not in the annual results. But the evidence of their consequences is accumulating.

Valencia: a €20 billion banking sector exposure, revealed in real time when the flood arrived.

Noto: a $7.5 to $17.6 billion economic damage event, occurring on a public holiday when monitoring was reduced and response capacity was compressed.

Canada: a record $8.5 billion insured loss year, with underwriting losses in personal property, reinsurance pricing increases of 25 to 70%, and a national insurance bureau warning that parts of Canada could become uninsurable within a decade.

These are not far-future projections.
They are today's operating environment.

> RESILIENCE USED TO MEAN HAVING A GOOD PLAN. TODAY IT MEANS HAVING INFRASTRUCTURE THAT CAN OPERATE WHEN THE PLAN RUNS OUT.

The NGFS revised its long-term damage estimates upward by a factor of three in November 2024, estimating economic losses of 15% of global GDP by 2050 under 2 degrees of warming. Six regulators across five jurisdictions have converged on mandatory operational integration of climate risk. Moody's RMS is classifying individual structures from satellite imagery within hours of a wildfire containment. Parametric insurance is paying out within days of a qualifying rainfall event in the Pacific Islands.

Financial institutions that build this infrastructure now will carry an advantage that cannot be assembled quickly during a crisis. The work begins with governed data foundations, continues through predictive systems that can operate continuously, and depends on response architecture that people trust enough to use. Munich Re's long-running natural disaster loss data shows the direction clearly. The harder question is whether financial institutions will have the sensing, coordination, and decision systems ready before the next major disruption tests them.

FULCRUM DIGITAL FD.

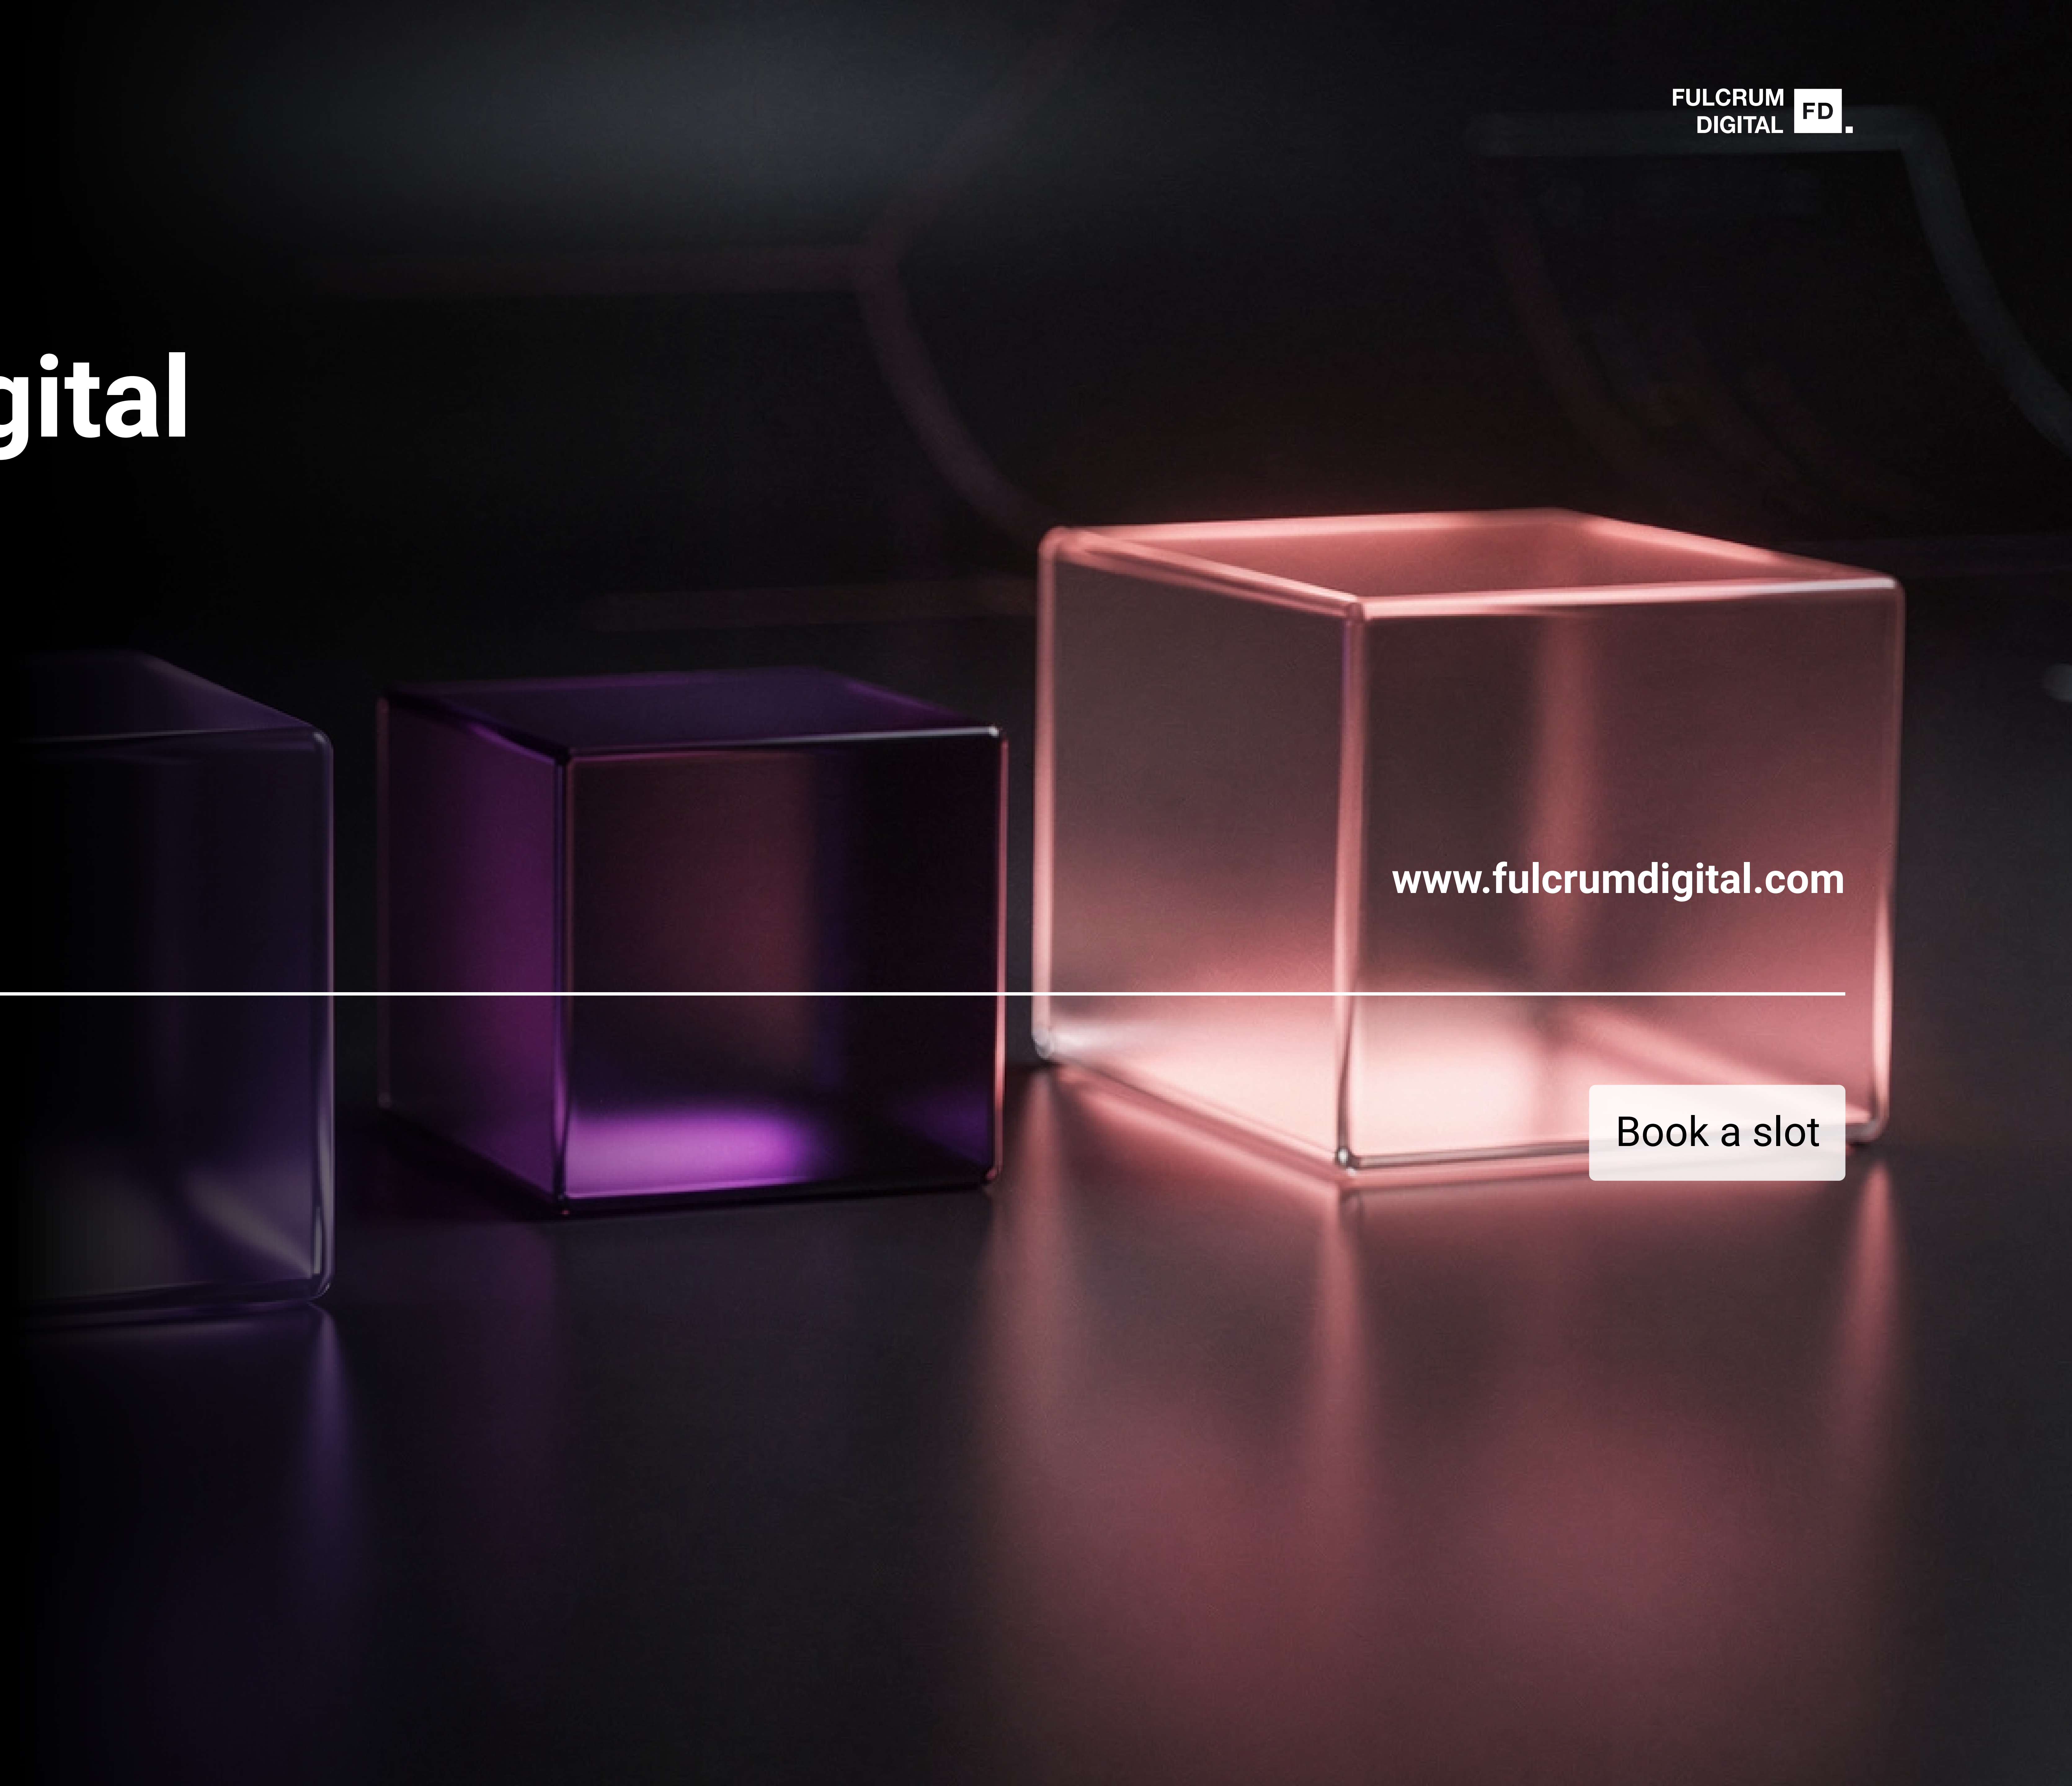

About Fulcrum Digital

Fulcrum Digital builds enterprise AI from architecture through production deployment and ongoing operations. With 1,500+ specialists, 4,500+ completed engagements, and production AI systems deployed across financial services, insurance, healthcare, and logistics since 1999, Fulcrum operates at the intersection of engineering depth, regulated-environment experience, and production discipline.

We build on your platform or ours.
You own everything we build.

www.fulcrumdigital.com

SCHEDULE A WORKING SESSION

For institutions exploring the architecture of operational intelligence, whether beginning with a data foundation assessment, extending existing AI capability toward response integration, or working through governance architecture for AI-assisted climate response, we welcome the conversation.

Book a slot