IROKO TECHNOLOGIES

WHITE PAPER
May 2026

Operational Topology

Understanding How Risk Actually Propagates Through Modern Enterprises

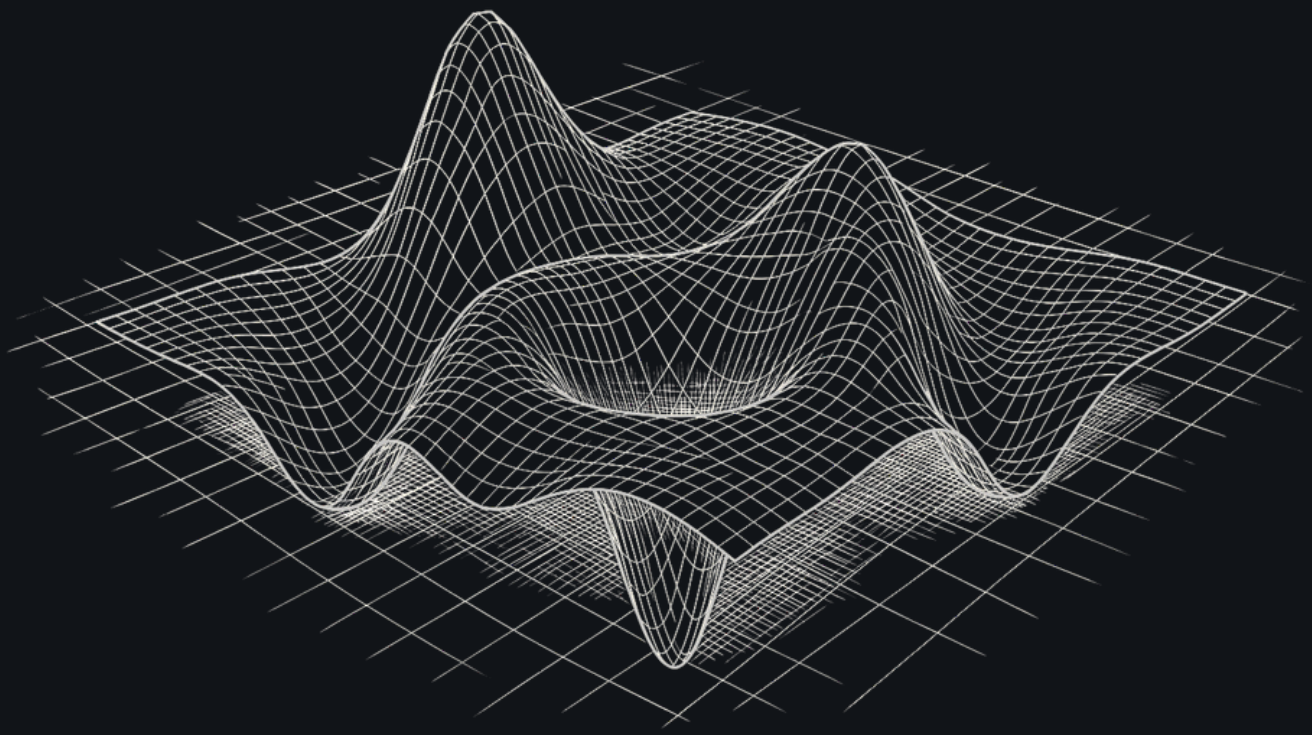

> *Operational risk is not only a property of components. In a networked enterprise, it is also a property of structure."*

JAMES HARDY

Co-Founder, Iroko Technologies
Former Global Head of Operational Resilience, State Street Corporation

IROKO
TECHNOLOGIES

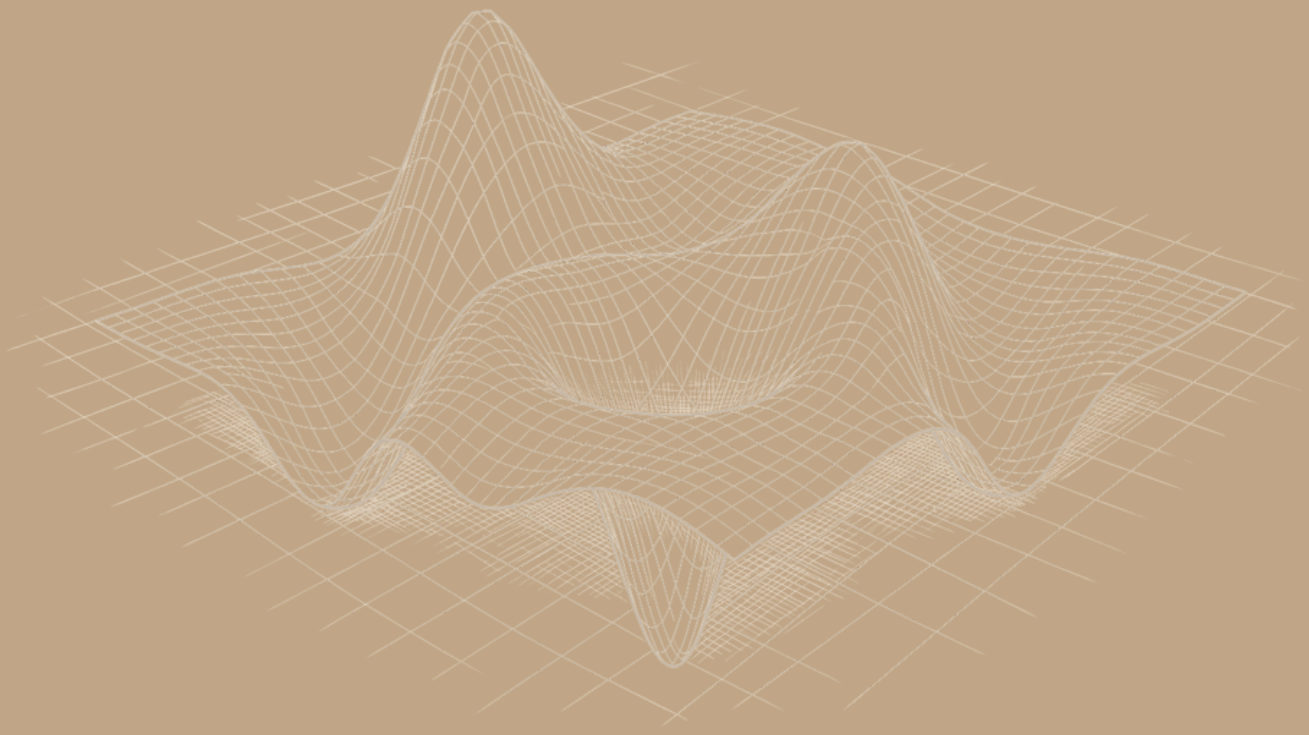

Executive Summary

Operational risk frameworks have done important work. They create accountability, support regulatory oversight, and give boards a structured view of non-financial exposure. But as enterprises have become more digital, interconnected, and dependent on shared platforms, those frameworks reveal an important limitation: they classify risk better than they represent how risk actually moves.

Most operational risk frameworks decompose exposure into domains such as technology, facilities, vendors, cyber, and process control. Those domain assessments are then aggregated into an overall view of risk. That approach is useful, but incomplete. Operational risk does not respect organizational or risk-domain boundaries. In modern enterprises, some of the most consequential exposures arise not from a single weak component, but from the structure of dependencies through which services are delivered.

Resilience disciplines have long recognized the importance of understanding how services depend on people, processes, technology, third parties, infrastructure, and locations. But that dependency view has only partially entered operational risk management in a systematic way. The result is a gap: operational risk often lacks the structural foundation it needs, while resilience planning is forced to build on an incomplete risk base.

This paper argues that operational risk management needs a structural lens. That lens is operational topology: the dependency network through which an organization actually delivers services and through which exposure builds, moves, and concentrates. Operational topology makes visible dependency importance, concentration, substitutability constraints, inherited exposure, and propagation paths that no individual domain assessment can reveal on its own.

The argument is not for replacing existing operational risk frameworks. It is for extending them. Domain-based assessments still matter, but on their own they can understate exposure because they do not fully represent the dependencies through which risk is inherited across the operating model.

A central question follows: if operational risk is shaped not only by the condition of individual components, but also by the way those components are connected, should frameworks explicitly measure that structure instead of merely documenting it?

The Flaw Built into the Foundation

Operational risk management emerged from a legitimate need: to give organizations and regulators a disciplined way to think about non-financial risks that can impair business performance. The frameworks that followed from the late 1990s onward — Basel II's operational risk capital requirements, COSO enterprise risk management, and the internal frameworks built on top of them — brought real rigor to risk oversight.

But many of the frameworks still used today were shaped in a less interconnected operating environment. In the early 2000s, enterprise technology was relatively modular. Systems were more siloed. Integration was limited. A failure in one area was less likely to propagate into another. Taxonomy-based risk management made sense in that environment because the systems it assessed were themselves relatively independent.

That is no longer the environment most firms operate in.

Today's enterprise operates as an interconnected system. A single customer-facing service may depend on dozens of applications, shared infrastructure platforms, multiple cloud providers, third-party software vendors, telecommunications networks, physical locations, and utility services. Those dependencies do not form a neat chain. They form a network — and in networks, failures propagate in ways that are difficult to infer from the perspective of any individual node.

The persistence of siloed risk assessment in this environment is not irrational. It reflects real advantages: accountability, specialist expertise, regulatory alignment, and cognitive manageability. But it also creates a systematic blind spot. When organizations assess risk domain by domain, they measure individual nodes while under-measuring the structure that connects them. In an interconnected system, a material share of operational exposure arises not from the condition of components alone, but from the architecture of dependency between them.

Why This Matters More Now

The argument for structural risk analysis is not new in academic circles. What is new is the operational reality that makes it urgent.

Thirty years ago, enterprise systems were relatively modular. A technology failure was largely contained within a technology boundary.

A facilities failure rarely cascaded into a service outage. The idea of a single vendor simultaneously disabling millions of endpoints globally would have seemed implausible.

Several converging forces have changed that calculus.

> Cloud Concentration

The migration to cloud infrastructure has delivered enormous efficiency gains while simultaneously concentrating critical dependencies onto a small number of hyperscale providers. A regional outage at a major cloud provider now disrupts thousands of organizations across unrelated industries at once. That concentration was not created by any individual organization's risk decision — it emerged from the aggregate of individually rational choices that produced a structurally fragile system.

> Software Supply Chain Depth

Modern applications are built on layered dependencies: open-source libraries, commercial software components, security agents, observability tools, and deployment frameworks. Each layer introduces dependencies that many application risk assessments do not surface explicitly. The attack surface — and the fragility surface — extends far below the applications that organizations formally assess.

> Platform and Ecosystem Dependencies

Enterprises increasingly rely on platform providers that sit beneath multiple internal systems simultaneously. A single integration platform, security tool, or identity provider may underpin dozens of applications. When such a platform fails, it is a structural event affecting every component it depends on, not just a technology risk event affecting one system.

> AI Infrastructure

The rapid integration of AI capabilities into operational processes is introducing a new layer of dependency into the enterprise. Organizations are beginning to rely on model providers, inference platforms, orchestration frameworks, and AI-enabled software components in ways that may sit beneath multiple services simultaneously. The structural risk lies not only in the failure of any one AI capability, but in the concentration that forms when many operational processes depend on the same underlying model, platform, or inference layer. That exposure is emerging faster than most risk frameworks are adapting.

> Human Capital Concentration

As organizations have become leaner and more specialized, human capital concentration has grown quietly alongside technological concentration. Technology is therefore not the only domain where structural concentration has intensified. Operational processes increasingly depend on individuals and teams who carry critical institutional knowledge, hold specialist skills, or occupy pivotal coordination roles across multiple services. A key person or team dependency is structurally similar to a single-vendor dependency: a node whose unavailability can propagate disruption across every process and service that relies on it.

The cumulative effect is clear: interconnectedness is no longer a secondary characteristic of the enterprise. It is the operating condition under which modern services are delivered.

How Operational Systems Actually Fail

The Cascade Problem

Operational services are delivered through layered dependencies. A realistic service stack looks something like this: a customer-facing service depends on a business process, which depends on application software, which depends on infrastructure, which depends on a data center, which depends on physical utilities. Each layer is typically managed by a different team and assessed under a different risk domain.

When a failure occurs, it travels through these layers. A cooling failure becomes infrastructure instability, which becomes application failure, which becomes customer impact. At each step, the failure crosses from one risk domain into another. A risk function assessing facilities, technology, and applications independently might find all three well-controlled — and still miss the cascade entirely. This is not a hypothetical concern. It is a recurring mechanism in many major operational failures.

Three Events That Illustrate the Gap

In July 2024, a faulty CrowdStrike Falcon content update caused Windows systems running the agent to crash and, in many cases, enter repeated restart states. Microsoft documented widespread endpoint failures and recovery actions, while CrowdStrike's own incident reporting identified the update as the direct technical cause.[1] The scale of disruption, however, was not explained by the defect alone. It was explained by structure: a fault in a widely shared operational dependency propagated simultaneously across large numbers of organizations.

[1]CrowdStrike, "Falcon Content Update Preliminary Post Incident Report," July 24, 2024. Available at: https://www.crowdstrike.com/en-us/blog/falcon-content-update-preliminary-post-incident-report/

CrowdStrike occupied a structurally critical position in many firms' operational environments: it operated with deep system access, was highly trusted, and was deployed broadly across enterprise endpoints. A conventional assessment of cyber risk or technology vendor risk, conducted within discrete domains, would have been unlikely to surface that structural exposure in full. The question that mattered was topological: how many critical systems shared the same dependency, and what concentration of exposure did that create across the operating model? CrowdStrike's own post-incident reporting identifies the affected Windows sensor/content update path, supporting the point that the event propagated through a common dependency layer.

The 2012 Knight Capital incident illustrates a different version of the same structural problem. A software deployment failure triggered unintended trading behavior, and within 45 minutes the firm had generated losses of approximately $440 million.[2] The event cut across technology, process, and operational control domains simultaneously, and it propagated through the system faster than effective human intervention could occur. What it exposed was not just a control failure at a single point, but a lack of structural visibility across the system as a whole.[3]

Major cloud outages reflect the same lesson at ecosystem scale. The December 2021 AWS US-EAST-1 disruption showed how disruption in a single cloud region can propagate simultaneously across many organizations that share the same provider-region dependency. The impact did not arise primarily from weaknesses in any one firm's internal controls. It arose from structural concentration in a common operational dependency. Firms may have assessed AWS as a technology or vendor risk but far fewer had measured how many critical services depended on the same provider, region, or architectural pattern.

The Pattern

In practice, the events that hurt most are often not the ones that begin in the obviously riskiest component. They are the ones that move through a dependency the organization relies on heavily and cannot easily route around. In those cases, the damage comes from the position that dependency occupies in the operating model: its importance, the concentration around it, the lack of credible substitutes, the breadth of its blast radius, or several of those at once.

[2]Knight Capital Group, Inc., "Knight Capital Group Provides Update Regarding August 1st Disruption to Routing in NYSE-Listed Securities," press release filed with the SEC, August 2, 2012. Available at: https://www.sec.gov/Archives/edgar/data/1060749/000119312512332176/d391111dex991.htm

[3]U.S. Securities and Exchange Commission, "SEC Charges Knight Capital With Violations of Market Access Rule," SEC Press Release 2013-222, October 16, 2013. Available at: https://www.sec.gov/newsroom/press-releases/2013-222

The risk is not simply a property of the component itself. It's also a property of the network in which it sits.

Siloed risk assessment is well designed to assess component-level risk but is poorly designed to assess structural risk. In increasingly interconnected operating environments, that gap is becoming more consequential.

Operational Topology: A Complementary Lens

Operational topology is the network of dependencies through which an organization delivers services and through which operational exposure is created, concentrated, and inherited. It treats the enterprise as a connected system rather than as a set of separately assessed domains.
The idea that network structure shapes systemic risk is well established in academic work. Haldane and May's work on banking ecosystems showed how interconnected financial networks exhibit a robust-yet-fragile property. They found banking ecosystems are stable under ordinary conditions, but capable of rapid systemic failure once a critical threshold is crossed.[4] Economists including Acemoglu, Ozdaglar, and Tahbaz-Salehi developed that logic further, showing formally how network topology, not just the condition of individual nodes, determines whether localized shocks remain contained or cascade across connected systems. [5]

Both bodies of work were developed in the context of financial networks and interbank contagion. The concern here is different: not how financial distress moves between institutions, but how operational exposure is created, concentrated, and inherited through the dependency structure of a single enterprise's operating model. The structural logic is related, even if the application is different.

The network has a specific anatomy. Nodes represent operational assets — services, processes, applications, infrastructure, locations, vendors, and the people and teams that operate them. Edges represent the dependency relationships between them. And the structure of that graph shapes where exposure accumulates, where failures spread, and where local weaknesses become systemic rather than contained.

By operational topology, I mean something specific: the map of how things depend on each other, and what that structure implies for operational exposure. What makes it analytically useful is not the map alone, but what the shape of that map reveals.

[4]Andrew G. Haldane and Robert M. May, "Systemic risk in banking ecosystems," Nature 469, no. 7330 (2011): 351–355. Available at: https://doi.org/10.1038/nature09659

[5]Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Systemic Risk and Stability in Financial Networks," American Economic Review 105, no. 2 (2015): 564–608. Available at: https://doi.org/10.1257/aer.20130456

Risk does not sit evenly across the network. It accumulates in some places more than others, and that pattern helps reveal forms of structural exposure that no assessment of individual components can show on its own.

> *"Operational topology does not just ask whether a component is risky. It asks what the organization is exposed to because of the way that component sits within the wider system."*

Both questions matter, but they answer different management problems. A component can have a low intrinsic risk rating and still represent a critical structural vulnerability because of where it sits within the dependency network.

Intrinsic Risk and Structural Risk

Intrinsic risk is the risk of a component considered in isolation: its reliability, security posture, control environment, and vulnerability to failure. Depending on the basis of assessment, that may reflect either an inherent or a residual view of risk. The distinction in this paper is not between inherent and residual risk, but between isolated component risk and exposure inherited through operational dependencies.

Structural risk is the systemic vulnerability created by that component's position within the dependency network: how strongly failure or disruption propagates through it, how many services rely on it, and whether credible alternative paths exist.

Domain-based frameworks are well suited to assessing the risk of individual components. Operational topology helps reveal the structural exposure created by the dependency network around them.

This distinction matters because a component can score well on intrinsic risk — robust controls, high availability, excellent vendor management — and still represent one of the organization's most significant structural vulnerabilities. The CrowdStrike agent was, from an intrinsic risk perspective, a mature and well-regarded security product. What domain-based assessment was less able to capture was the structural exposure created by its position across many critical dependency paths.

Traditional frameworks are not wrong to focus on intrinsic risk. But in a networked system, intrinsic risk without structural context produces an incomplete picture of exposure. Only by combining the two can organizations see both the condition of individual components and the way risk behaves across the system.

Intrinsic Risk and Inherited Exposure

Operational components carry not only intrinsic risk — the risk arising from their own control posture and operating condition — but also inherited exposure: the exposure transmitted upward through the dependencies that support them.

A customer-facing service can be well designed, well monitored, and well controlled. But if it depends on an application with weaker controls, which in turn depends on infrastructure with a single point of failure in a less well controlled facility, the service still inherits exposure from the dependency layers beneath it — exposure that would not be visible from assessing the service on its own.

This has a direct practical implication: two services with identical intrinsic risk profiles can have radically different effective exposures depending on the topology of their dependencies. The service that converges on a shared infrastructure component at a point of concentration inherits a structural exposure that the service with independent, redundant dependencies does not. No domain-based assessment of either service individually will surface this difference. A topological view is what reveals it.

The Importance of Dependencies

Not all dependencies carry equal operational significance. A critical payment service may rely on an authentication platform as a mandatory control point, while another service uses the same platform only for optional functionality. The structural significance of these two dependencies is entirely different.

Operational topology must therefore account not only for the existence of dependencies, but for their importance within the service path. A dependency that is essential to service delivery transmits weakness more strongly than one that is peripheral or easily substituted.

Effective exposure is therefore not simply the sum of a component's dependencies. It is shaped by the importance of each one — and a single mandatory dependency on a vulnerable component can outweigh a dozen peripheral ones.

THE ANATOMY OF EFFECTIVE EXPOSURE

The effective exposure of any operational component arises from two sources: its intrinsic risk and the exposure it inherits through the dependencies beneath it. Because not all dependencies matter equally, inherited exposure should be understood as importance-weighted rather than uniformly distributed across the operating model.

Effective Exposure = Intrinsic Risk + Inherited Exposure

Inherited Exposure = the combined effect of dependency risk and dependency importance

This is a conceptual expression rather than a fixed mathematical specification. In practice, organizations may calibrate inherited exposure differently depending on their operating model, including adjustments for redundancy, recovery characteristics, path depth, or non-linear propagation effects. The essential point is simpler: effective exposure cannot be understood from the condition of a component alone. It also depends on the structure through which exposure is inherited.

Topology analysis is therefore not a substitute for control assessment. It is a way of understanding how control weakness, concentration, and dependency structure interact to produce effective exposure.

In highly interconnected systems, inherited exposure may equal or exceed intrinsic risk. A component that appears well controlled in isolation may still carry substantial effective exposure if it sits at the convergence of high-importance dependency paths.

Visualizing the Operational Risk Landscape

Elevation reflects effective exposure. A node rises in the landscape when its own intrinsic risk is compounded by inherited exposure from the dependencies beneath it. Nodes with lower intrinsic risk, stronger supporting dependencies, or lower-importance relationships sit lower in the landscape.

The highest peaks form where multiple vulnerable or high-importance dependency paths converge. These are the parts of the operating model where exposure has accumulated most sharply. The steepest slopes mark the pathways along which inherited exposure propagates most strongly.

But elevation is only one dimension of the risk landscape. A node also has a blast radius: the extent to which disruption would spread if that node failed.

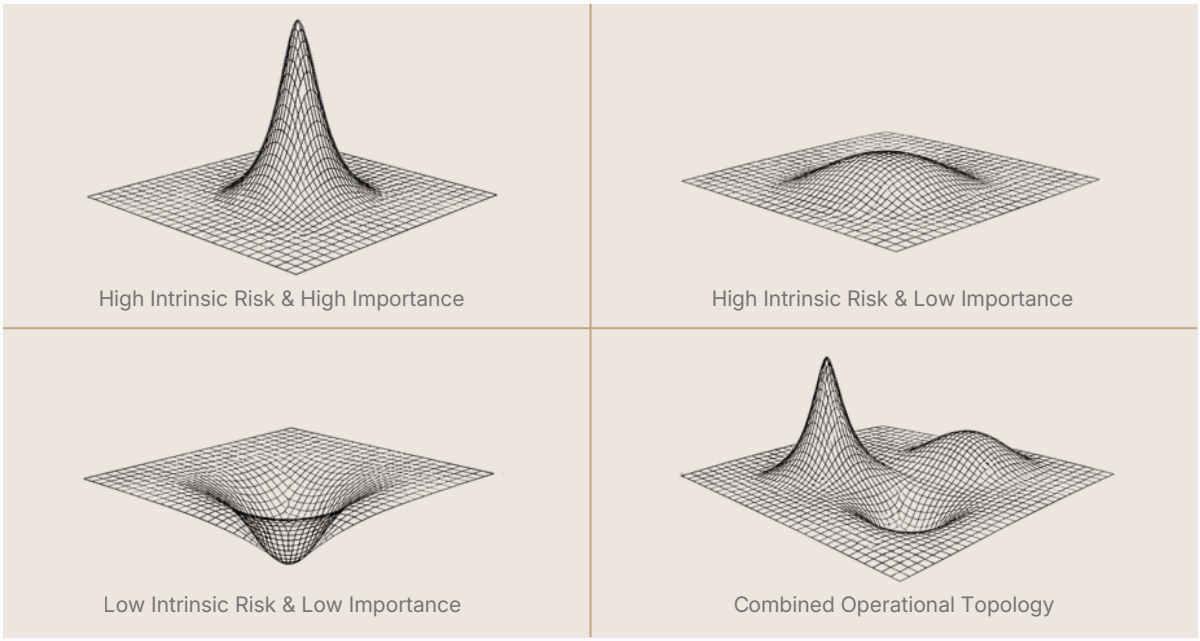

Figure 1: Operational Risk Topology. Each node's elevation reflects its effective exposure — the combination of intrinsic risk and importance-weighted inherited exposure flowing through its dependencies. Steeper slopes indicate critical operational pathways along which risk propagates more strongly. Nodes appearing as valleys carry lower effective exposure or lower-importance relationships.

Elevation and blast radius are distinct. A node may sit at moderate height in the landscape yet underpin many critical services. When it fails, disruption reaches wide — not because the node was highly exposed, but because of where it sat within the network. This was the CrowdStrike condition: moderate elevation, exceptional blast radius.

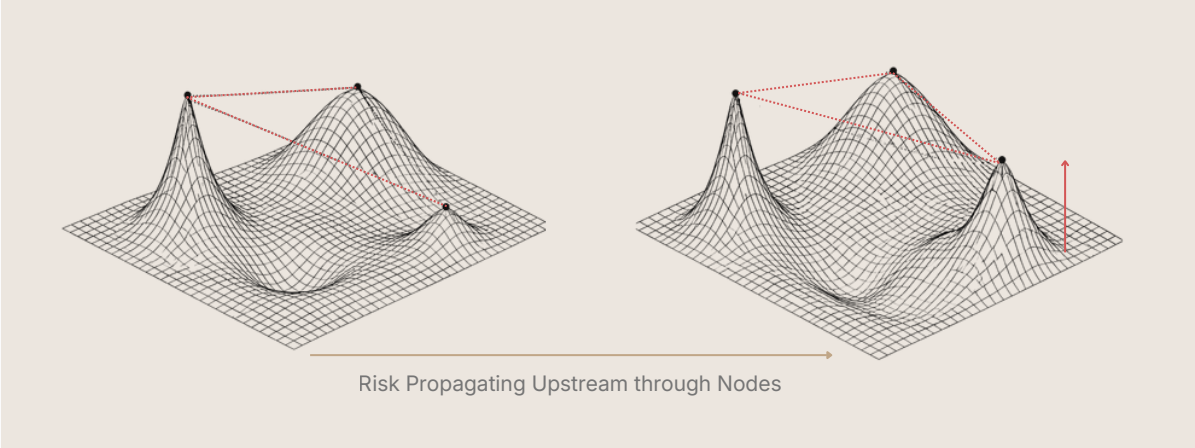

Figure 2: Risk Propagation Across the Operational Topology. The blast radius of disruption reflects structural connectivity and dependency importance — not the elevation of the originating node. A node of moderate systemic exposure can produce wide-reaching disruption if it sits at the convergence of many critical dependency paths.

Centrality — how connected a node is — can suggest where to look, but it is not a reliable guide to effective exposure. The most dangerous nodes are often not the most connected ones. They are the ones sitting on paths the enterprise cannot route around.

Taking the Counterarguments Seriously

Any proposal to augment established frameworks must contend with real objections. Four deserve substantive engagement.

Objection 1: Business Continuity Already Exists

Business continuity and disaster recovery programmes do capture some dependency dynamics. Good programmes identify critical processes, supporting teams, technology, facilities, and third parties, and use scenario testing to assess whether services can be recovered within tolerance. The issue is not that continuity disciplines overlook dependencies. It is that they map them for a different reason, and that difference is more important than it may first seem.

BCP uses dependency mapping to support recovery. It starts with a service disruption and works outward: what people, systems, facilities, suppliers, and workarounds are needed to restore the service within tolerance? Operational topology starts in a different place. It looks at the structure of the operating model itself and asks where risk is already concentrated, where exposure is inherited, and where failures would be most likely to spread.

Those are related questions, but they are not the same. Continuity analysis is usually tied to a defined disruption scenario. Topology analysis does not need that starting point. It can surface concentration, weak substitutability, and structurally important dependencies before a particular failure has been specified.

This distinction has a practical consequence that goes beyond resilience. Structural exposure is an operational risk management problem, not only a continuity planning problem. A concentration of critical services on a single vendor or platform is a risk condition that should appear in operational risk reporting, inform risk appetite discussions, and drive risk treatment decisions — regardless of whether a recovery plan exists for it. In most organizations today, it does neither. Dependency mapping typically lives in resilience functions and rarely makes its way into operational risk frameworks in a systematic way. The result is that operational risk management often remains only partially sighted to a category of exposure that its own frameworks are intended to capture.

The CrowdStrike event did not match any BCP scenario for most affected organizations — not because the scenario was unimaginable in principle, but because the structural exposure that made it devastating had never been measured as a risk management matter. In many organizations, no explicit risk treatment had been applied, no risk appetite threshold had been framed in structural terms, and no board reporting had highlighted the concentration. The exposure remained largely invisible to operational risk management until it became a crisis.

Operational topology does not replace business continuity. It provides the structural foundation that operational risk management is currently missing — and from which more realistic, more targeted continuity planning can then be built.

Objection 2: Mapping Dependencies at Scale is Impractical

This is the strongest objection and should not be underestimated. Large organizations operate thousands of applications, hundreds of vendors, and dozens of physical locations. Maintaining an accurate, complete map of all dependencies is a significant undertaking.

The answer is prioritized sufficiency, not exhaustive completeness. Topology modeling does not require mapping every dependency in the enterprise from the outset. Start with the services the organization cannot afford to lose — those are known before any mapping begins. The dependencies that matter most will surface through that process: the shared nodes that appear repeatedly, the paths with no alternative, the components buried deep beneath multiple critical services.

Mapping reveals what matters; it does not require you to know it in advance. A risk-stratified topology built iteratively around critical services provides most of the analytical value at a fraction of the cost of full enterprise mapping. Modern integration with CMDB systems and vendor risk platforms is also making this progressively more tractable. The data challenge is real but diminishing.

Objection 3: Governance Requires Classification

This objection is correct, and topology modeling does not contest it. Accountability structures require that someone owns each risk domain. Regulatory oversight is organized around risk categories. Specialist expertise is developed within domains. None of this should change.

The claim is not that classification should be replaced. It is that classification alone produces an incomplete picture of operational risk in a networked system. Who is responsible for managing a risk is an entirely separate question from how that risk

behaves within the system as a whole. Organizations can maintain full accountability structures for each domain while simultaneously developing a cross-domain view of structural vulnerability. These perspectives do different work.

Objection 4: Quantitative Risk Frameworks Already Model This

Practitioners familiar with Factor Analysis of Information Risk (FAIR) will note that quantitative operational risk frameworks already decompose risk into components and can, in principle, model dependency-driven loss cascades. That is true, and FAIR-based programmes represent an advance over purely qualitative approaches in that respect.

However, the distinction is one of starting point and scope. FAIR is scenario-driven: it quantifies expected loss arising from a defined threat event acting upon a defined asset. It is well suited to answering, "How bad could this be?" for a known risk. Operational topology, on the other hand, is structure-driven: it analyzes the dependency architecture of the operating model and makes concentration, propagation paths, and inherited exposure visible before any specific threat scenario is defined. It is therefore designed to answer, "Where are the structural vulnerabilities we have not yet thought to ask about?" Threat events and likelihood therefore remain important, but they are not the analytical starting point. Topology instead helps determine where those lenses should be applied by identifying the parts of the operating model where disruption would propagate most severely, where concentration is greatest, and where substitutability is weakest.

The two approaches are thus complementary: topology helps surface structurally material exposures, and scenario-based quantification can then estimate their loss implications.

Where Regulations is Already Heading

The direction of travel in operational resilience regulation already supports the topology argument, even if the explicit language of topology is not yet there. The most developed formal regimes are concentrated in financial services, particularly in the UK and EU. Frameworks introduced by the PRA and FCA in the UK, DORA in the EU, and related principles issued by the Basel Committee all begin from a similar logic: identify important business services or critical operations, map the dependencies and interconnections that support them, define tolerances, and demonstrate resilience under disruption.[6][7][8][9]

[6]Financial Conduct Authority, "PS21/3: Building operational resilience," March 29, 2021. Available at: https://www.fca.org.uk/publications/policy-statements/ps21-3-building-operational-resilience

[7]Prudential Regulation Authority, "SS1/21: Operational resilience – Impact tolerances for important business services," March 2021; updated March 2022 PDF. Available at: https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2021/ss121-march-22.pdf]

This is topology-adjacent thinking, but it is not yet topology measurement embedded in operational risk management. Regulators are increasingly asking firms to identify important services, map dependencies, explain substitutability constraints, and demonstrate they can remain within impact tolerances under realistic stress. Those are structurally informed questions. But they are still asked mainly through a resilience and third-party risk frame focused on continuity, tolerance, recovery, and concentration management rather than as an explicit requirement to identify, measure, and manage structural exposure within operational risk frameworks themselves. The result is that a firm may be able to satisfy a resilience scenario or outsourcing expectation while structural concentration remains only partially reflected in risk registers, risk appetite, or treatment decisions. That is the gap topology-aware methods are intended to close.

The trajectory matters as much as the current requirements. Financial services regulators have moved on from simply asking firms to document risks. The logical next step would be requiring firms to show they understand how structural concentration and inherited exposure affect those tolerances. It is the same question, asked more precisely. Organizations that have built a structural view of their operating model before that requirement is formalized will be better positioned to respond and better positioned to engage credibly with regulators and boards in the meantime.

The first-mover advantage here is real. Regulatory requirements in this space tend to be calibrated to what leading firms have already demonstrated is achievable. Organizations that develop topology capability now are not only managing risk more effectively — they are shaping the practical standard against which future requirements are likely to be set.

Toward a Topology-Aware Risk Programme

Moving from domain-based to topology-aware risk management does not require organizations to rebuild their existing programmes. It requires adding a structural layer to operational risk itself: one that identifies material dependency paths, measures inherited exposure, and provides a common analytical foundation for continuity planning, resilience testing, and domain risk assessment.

This implies a meaningful shift from how many organizations still operate today. Dependency mapping often sits primarily within resilience functions, where it is used to support recovery planning and scenario testing.

[8]Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA). Available at: https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng

[9]Basel Committee on Banking Supervision, "Principles for operational resilience," March 2021. Available at: https://www.bis.org/bcbs/publ/d516.pdf

Operational risk functions may consume some of its outputs, but often only indirectly and inconsistently. A topology-aware programme reverses that relationship. The structural view of the operating model becomes a core risk management asset in its own right, with resilience planning drawing from it rather than developing it separately. In that model, operational risk establishes the structural foundation; resilience planning builds on it.

In practical terms, a topology-aware programme involves five recurring activities: identifying the services that matter most, mapping their material dependencies, understanding structural exposure and where it concentrates, maintaining that view as the operating model changes, and integrating structural insight with existing domain assessments.

Start With Services, Not Domains

Topology mapping should begin with a simple question: which services does the organization most need to protect? Services define what matters. Dependencies define what those services rely on. Risk domains describe how those dependencies are governed and managed. Starting with services ensures that topology analysis remains anchored in business outcomes rather than organizational structure. It also creates a common point of reference for operational risk and resilience teams, reducing the tendency for each to work from parallel but only loosely connected views of the operating model.

Map for Structural Materiality, Not Completeness

The purpose of initial topology mapping is not to build an exhaustive inventory of every operational relationship. It is to identify structurally material vulnerabilities: high-importance dependencies through which disruption would propagate most strongly, shared dependencies with the greatest concentration risk, and redundancy gaps where no credible alternative path exists. Mapping can begin with important services and their most material dependency paths, then deepen iteratively as further concentrations and hidden exposures are revealed. The value of topology analysis does not depend on completeness at the outset.

Surface Structural Risk Alongside Domain Risk

Executive and board reporting should present a structural view alongside conventional domain assessments. That view should make visible the most important dependency paths, the largest blast radii, the strongest concentrations, and the most consequential redundancy gaps. This allows leadership to see not only where risk sits within individual domains, but how it is capable of propagating across the operating model. For senior management, including COOs and CIOs, this also provides a practical management

lens: a living view of where the organization is most fragile and where investment in redundancy, substitutability, or architectural change is likely to reduce risk most materially.

Treat the Topology as a Living Asset

An accurate operational topology is not merely a resilience artefact produced for regulatory review and then allowed to age. It is a risk management asset that supports better scenario design, more targeted resilience investment, and more credible engagement with regulators and senior stakeholders. To remain useful, it must be maintained as the operating model changes: when applications are added or retired, vendors change, infrastructure is modified, or organizational dependencies shift. That maintenance should be embedded in the risk management process rather than handled as an isolated resilience exercise.

Integrate with Existing Domain Assessments

Topology analysis should both inform and be informed by domain-based assessments. A vendor risk review that identifies a supplier as high risk becomes more decision-useful when combined with topology analysis showing how many important services depend on that supplier and whether viable substitutes exist. A technology risk assessment highlighting an aging infrastructure component becomes more meaningful when the structural blast radius of its failure is visible. Neither view is sufficient on its own. Together, they provide a fuller picture of both the sources of operational weakness and the ways in which that weakness can propagate.

Conclusion

Domain-based operational risk frameworks have earned their place in governance. They create accountability, enable specialist expertise, and provide the structured oversight that regulators require. Operational risk management is more disciplined today because of them.

But operational systems have changed faster than the frameworks designed to assess them. In a connected, digitally integrated enterprise, a meaningful share of operational exposure arises not only from the condition of individual components, but from the structure of dependency between them. That is where concentration forms, substitutes become limited, dependency importance becomes clearer, and exposure is inherited upward through the operating model.

Operational topology makes that structural dimension visible. It helps organizations see which dependencies matter most, where risk is concentrated, where structural exposure is greatest, and where conventional domain assessments,

on their own, can understate exposure. It does not replace existing frameworks; it fills in something they often miss.

The more useful distinction is between assessing components one by one and understanding the structure that connects them. Organizations that can measure both will have a more accurate view of operational risk, a stronger basis for prioritization and investment, and a clearer way to identify vulnerabilities that no single domain can fully see in isolation.

In a networked enterprise, structure is not just background. It is part of the risk.

About the Author

James Hardy has held senior roles across technology, operations, and operational resilience in financial services including serving as Global Head of Operational Resilience. He is the Co-Founder of Iroko Technologies, the company behind OneScor, a platform for operational risk and resilience intelligence.

This paper reflects a practitioner perspective developed from building and running operational risk and resilience programmes across complex financial institutions, and from seeing first-hand the structural limitations of domain-based frameworks in highly interconnected operating environments.

irokotechnologies.com

IROKO TECHNOLOGIES | OneScor