IROKO TECHNOLOGIES

neo4j

Measuring Structure

Beyond dependency mapping:
how OneScor and Neo4j turn operational topology into defensible risk and resilience decisions

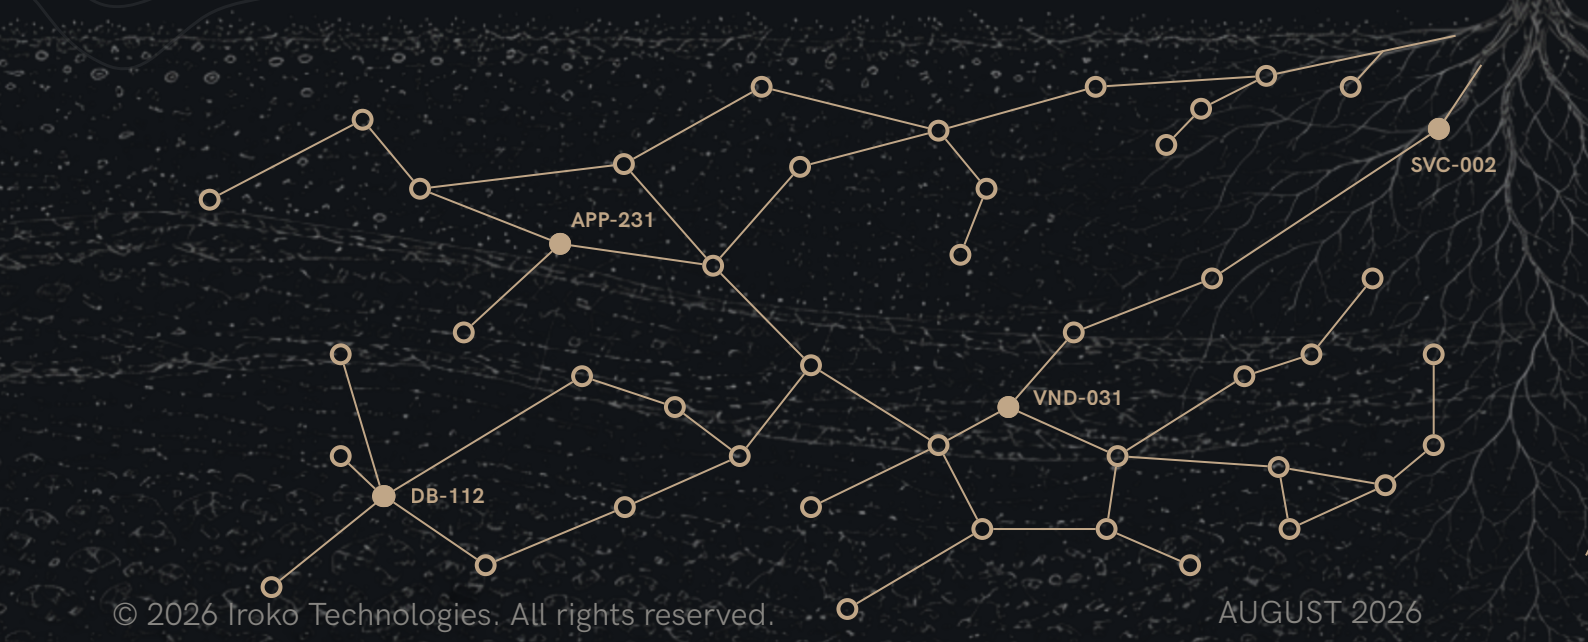

© 2026 Iroko Technologies. All rights reserved.

AUGUST 2026

01/ Executive Summary

> A dependency map shows where the connections are. Measurement shows how connections impact exposure, and which actions matter most.

Firms have spent years mapping the people, technology, processes, facilities, information and third parties that support important business services. That work matters: it reveals how the business operates and where disruption can travel. Regulators have reinforced the same expectation. Basel, the UK supervisory authorities and DORA require firms to understand the resources and dependencies behind important services and use that understanding in testing, recovery and third-party oversight. US agencies have taken a similar direction through operational-resilience and third-party-risk guidance.

But a connected view is not yet a risk measure. A map can show that several important services rely on the same component. It does not, by itself, establish how much exposure that component carries, whether apparent alternatives are genuinely independent, or which remediation would reduce the most consequential risk.

OneScor, Iroko Technologies' risk and resilience platform, applies governed operational risk and resilience logic to that connected structure. It is built on what we call the Operational Topology. Neo4j provides the graph foundation that makes traversal, reachability and path analysis direct and scalable at enterprise depth.

Together they turn dependency information into a decomposable view of effective exposure: a component's assessed condition combined with structural risk inherited through its dependencies. The result is not another diagram but a governed model for finding concentrations, challenging redundancy, prioritizing treatment, directing testing and explaining every decision. It can also be tested against incident experience and recalibrated as evidence improves. **The map is the input. The decision requires measurement.**

What is Operational Topology?

Operational topology is the connected structure through which a firm actually operates: its important services, processes, applications, infrastructure, third parties, information, locations and teams, and the dependencies that join them. An org chart shows how a firm is arranged; a risk register shows what it worries about. Operational topology shows what carries the business, and therefore what carries the risk.

A companion paper, *Operational Topology,* develops that argument. This paper is about measuring it.

© 2026 Iroko Technologies. All rights reserved.

02/

Mapping alone can’t answer the questions

Operational risk is typically assessed by domain and at the level of the individual component being evaluated. Technology risk assesses applications and infrastructure. Facilities and property teams assess sites and buildings. Country and geopolitical risk functions assess locations. Third-party risk assesses the external supply chain. These assessments may share a common GRC platform, enterprise taxonomy or rating framework, but their conclusions generally remain bounded by the component and domain under review. They are rarely combined through the structural relationships that connect the operating model.

The questions that matter most, however, cut across those domains: how exposed an important service really is, where that exposure comes from and which failures would cause it the greatest harm.

Consider a deliberately simple example. A firm consists of a single business line delivering one service through one supporting process. That process depends on one application, which is located in a single data center.

Each component has been assessed within its relevant domain. A Risk and Control Self-Assessment (RCSA) rates the process as *Low* residual risk, a technology risk assessment rates the application as *Moderate*, and a location assessment rates the data center as *High*. Each rating describes a different component within the scope and methodology of the assessment that produced it. Before these ratings are used in the calculation, they are translated through governed mappings onto a common exposure scale, while the original ratings, methodologies, and lineage are preserved.

FIGURE 1 WHAT DOES EACH ASSESSMENT SEE ON ITS OWN?

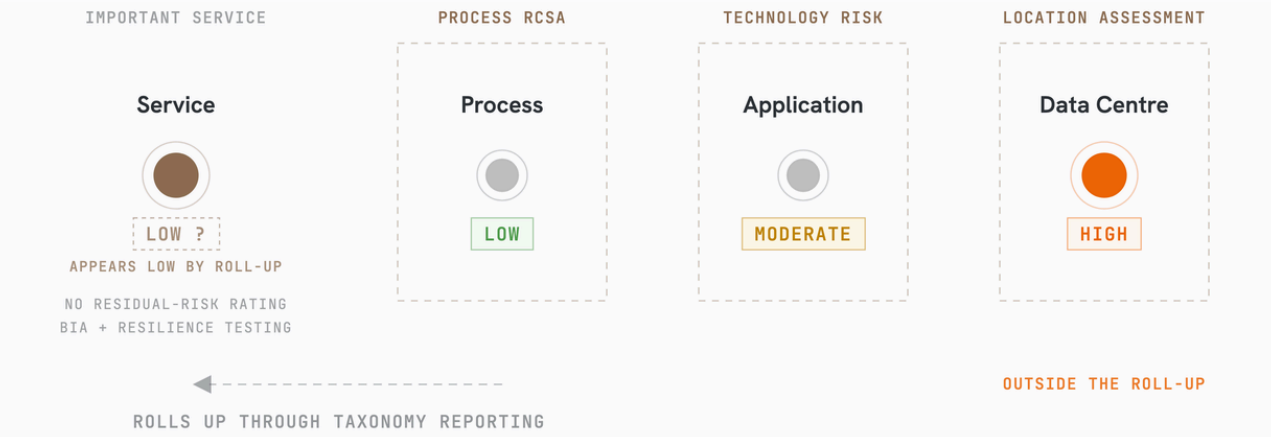

Figure 1. The example as each domain assesses it. Every rating is valid within its own scope, and the only supporting process RCSA is LOW, which can leave the impression that the service is similarly low risk.

© 2026 Iroko Technologies. All rights reserved.

The service itself may not have a conventional residual-risk assessment. Its Business Impact Analysis (BIA) establishes the consequences of disruption, while operational resilience activities assess its ability to remain within tolerance. In many firms, the process RCSA would otherwise roll up through organizational and risk-taxonomy reporting. In this simple example, the only supporting process RCSA is *Low*, potentially creating the impression that the service is similarly low risk.

The picture changes when the assessments are evaluated across the structural relationships between the components. **Relies on**, **depends on** and **hosted in** are not merely descriptive links. They are pathways through which exposure is inherited and disruption can propagate.

Evaluated across that path, the service's effective exposure resolves to *High*. The result is driven by the data-center location assessment, even though that exposure is not visible in the process RCSA or its conventional reporting roll-up.

FIGURE 2 HOW EXPOSED IS THE SERVICE, ONCE THE STRUCTURE IS CONSIDERED?

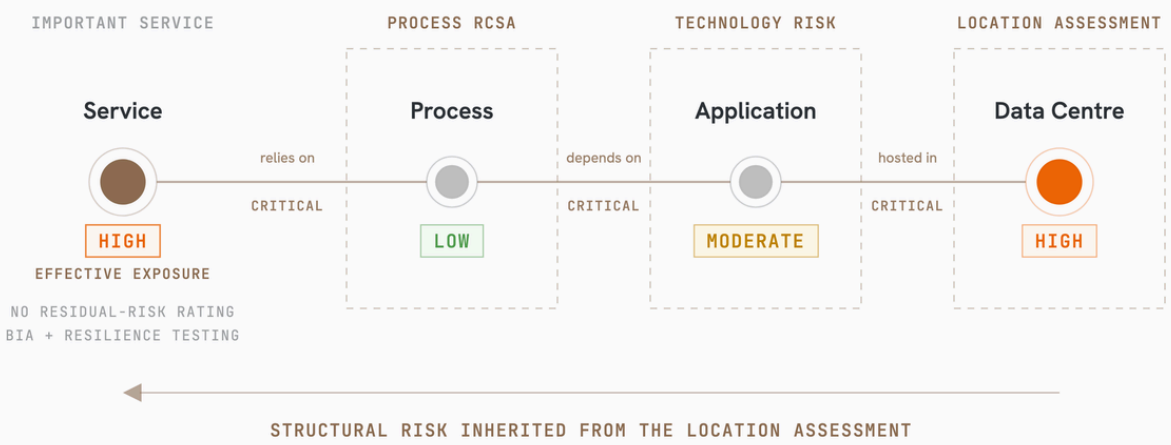

Figure 2. The same components, with the structural relationships drawn. Relies on, Depends on and Hosted in are relationships that risk traverses. Measured through them, the service's effective exposure may well be HIGH, inherited from the location assessment rather than from anything the roll-up reached.

> *The map exposed the dependencies. Measurement revealed a different outcome.*

In mature firms, everything in this example is likely to be documented somewhere. What is often missing is not awareness, but a repeatable, scalable and decomposable computation across the firm's operating structure. **Effective exposure is not solely a property of an individual component.** It is a function of both the assessed condition of the components and the structural paths through which their risk can propagate. **Mapping identifies those paths; computation determines their effect.**

Measuring exposure across domains changes both understanding and action. The service owner can see that the High data-center exposure materially drives the service's position, even though it sits outside the conventional process roll-up. The technology remediation backlog is no longer automatically the priority: further hardening the *Moderate* application may improve its local condition but leave the service's *High* rating while the location dependency remains the dominant source of exposure.

© 2026 Iroko Technologies. All rights reserved.

FIGURE 3 WHAT DOES EACH STEP MAKE VISIBLE?

01/ **Assessment View**

The process is rated LOW by RCSA, the application MODERATE by technology risk, and the data center HIGH by the location assessment. The service itself carries no residual-risk rating.

Valid assessments, each bounded by its own scope and methodology.

02/ **Network Analysis**

The service relies on the process, which depends on the application, which is located in the data center. The path crosses all three assessment scopes and reaches a HIGH-rated location.

A dependency that neither a single assessment nor the reporting roll-up is scoped to reach.

03/ **OneScor View**

The data center's structural risk is carried up the path into the service's effective exposure, which may well resolve to HIGH rather than the LOW suggested by the process RCSA.

A defensible priority: address the location dependency rather than further harden an application.

Figure 3. Each step makes something new visible, and each establishes something the step before it could not. The reader has just watched all three happen in the example.

The example is deliberately simple and follows exposure to a single service. In reality, an important service may depend on hundreds or thousands of processes, applications, infrastructure components, locations and third parties connected through overlapping, many-to-many paths. The same component may support multiple services, appear in several dependency chains and play a different role in each. Nor are all relationships equally important: a dependency may be critical to one service but readily substitutable for another, while a concentrated dependency with no viable alternative should not be treated like an optional connection.

Effective exposure is therefore calculated at every node in the topology, not only at the service level. Each component carries its own exposure, combining its assessed condition with the exposure inherited through its dependencies. That calculation must consider which components it depends on, the paths through which they are connected, and the nature, context and importance of each relationship. **At enterprise scale, the challenge is not simply to draw a larger map, but to maintain a standing, decomposable measure of exposure across the complete operating topology.**

FIGURE 4 FROM A SIMPLE PATH TO OPERATING REALITY

Figure 4. Operating reality is a dense, overlapping topology rather than a single chain. Every node carries its own effective exposure, shaped by its assessed condition, its dependencies and the nature and importance of the relationships connecting them.

© 2026 Iroko Technologies. All rights reserved.

The effective exposure carried by each component—combining its own assessed condition, where one exists, with the structural risk inherited through its dependencies—is a measure the map alone cannot produce.

03/ Effective exposure: condition and position

The final High rating in the example was not a direct assessment of the service or any single component. It was computed from the assessed condition of the components and the structure beneath the service.

Intrinsic risk is the risk attributable to a component based on its own assessed condition, independent of any exposure inherited through its dependencies. It may be represented by a residual-risk rating, an inherent-risk rating or another governed assessment result, depending on the methodology selected for that component.

Before different assessment results are combined, they are normalized to a common exposure scale used to calculate effective exposure. Effective exposure combines a component's intrinsic risk with the structural risk inherited through its dependencies and relationships.

FIGURE 5 WHAT IS EFFECTIVE EXPOSURE MADE OF?

INTRINSIC RISK	+	STRUCTURAL RISK	=	EFFECTIVE EXPOSURE
The component's own governed risk assessment, normalized to the common scale used by the model.		The risk inherited through its dependencies and relationships.		The governed combination of both.

Figure 5. The anatomy of effective exposure. Two services with the same intrinsic rating can carry different effective exposure because the structures beneath them differ.

Together, intrinsic and structural risk capture condition and position. Services are delivered through layers owned by different functions and assessed through different domains, but risk and disruption do not respect those boundaries: a location failure becomes an application failure, then a process and service failure resulting in customer impact.

04/ What defensible structural measurement requires

Producing a rating is easy. Producing one that a risk and resilience team can examine, compare, govern and defend requires commitments a map alone does not make. OneScor's structural measurement method is built around five commitments.

© 2026 Iroko Technologies. All rights reserved.

01/ **Importance lives on the relationship.**

A mandatory control point and an optional dependency should not transmit risk the same way. Importance therefore belongs on the relationship, not on the component.

02/ **Risk travels by a stated rule.**

The way risk propagates through the structure requires a rule that must be explicit, consistent and versioned because a result produced by an unstated rule cannot be examined or defended.

03/ **Redundancy counts only when it is real.**

An alternative should reduce exposure only when it is genuinely usable: independent, sufficiently capable and recoverable within the required time.

04/ **Every result decomposes.**

Every result should be examinable, on demand, back to the inputs, rules and dependency paths that produced it. Without that it is an assertion; with it, it is evidence.

05/ **The method is governed and validated.**

The scales, thresholds and rules that produce a rating can be based on expert judgment, approved, versioned and well governed but should be open to challenge based on empirical evidence.

Together, these commitments are what separate measuring the map from simply assembling it: they make the measurement reproducible, challengeable and capable of improvement. They are also why the computational model must treat connections as first-class, because importance is held on a relationship, propagation follows paths, redundancy requires an independent path and decomposition retains the traversed path.

05/ Why the computational model is a graph

> *"In organizations of this complexity, the most critical risks aren't in the assets; they're hiding in the connections between them."*
>
> Michael Down
> Global Head of Financial Services, Neo4j

© 2026 Iroko Technologies. All rights reserved.

Traditional data structures treat assets as static islands, forcing risk teams to manually "reconstruct" paths every time they need to answer a question. This is not only slow; it is brittle. It leads to fragmented views where the truth is obscured by the difficulty of connecting silos. To move from static mapping to dynamic risk management, **we need a model where connections aren't afterthoughts; they are first-class citizens.**

By representing operational topology that mimics the real world, we break the cycle of "collecting data" and start "connecting meaning." This is the fundamental shift required for managing risk in today's organizations:

01/ **From blind spots to blast radius:**

In a graph, complex questions, such as "What is our total exposure if this specific broker fails?" or "Do we have a genuinely independent alternative?" move from multi-day manual exercises to instant, accurate traversals. We are not just mapping; we are navigating risk.

02/ **From fragmented silos to shared institutional memory:**

Because the graph serves as a single, governed substrate, it eliminates the "context gap" in which every team rebuilds business logic from scratch. It creates a cumulative, authoritative view of how the organization actually operates.

03/ **From passive collection to defensible measurement:**

When connections are treated as first-class, we enable a "computable organization." This transforms complexity from a liability into a clear, actionable view of exposure that can be audited, challenged, and trusted.

This architectural shift provides the foundation, but requires a platform to apply it. The practical question is how to run this measurement continuously across thousands of components and dependencies, under governance that an auditor or regulator can inspect.

06/ OneScor computes what the map means

OneScor measures the exposure carried by the structure of the firm, and turns that measure into action: where the risk sits, why it is there, and which fix matters most. **It is not a mapping tool with ratings attached but the five commitments turned into software, operating against the real model of the firm constructed on Neo4j.**

Its capabilities build on one another: mapping is the foundation; structural exposure is the step this paper has argued the map alone cannot compute, and four further capabilities put the score to work.

© 2026 Iroko Technologies. All rights reserved.

OneScor

Map FOUNDATION
- Unified asset inventory
- Dependency mapping
- Critical path identification

Score FOUNDATION
- Factor-based assessments
- Dependency-aware propagation
- Fully traceable scores

Impact
- Cascading impact and blast radius
- Time-based impact modeling
- Resilience gap analysis

Simulate
- Full-cycle simulation
- Operational playbook generation
- Lessons learned and evidence

Optimize
- Structural analytics
- Prescriptive intelligence
- Scenario prioritization

Reason
- Traceable answers
- Workflow automation
- Credibility-tiered research

> *"OneScor is the platform I wish I had when the board was asking me 'Are our critical services actually resilient?' and 'what is the most effective way to reduce risk?'"*
>
> James Hardy
> Co-Founder, Iroko Technologies
> Former Global Head of Operational Resilience, State Street Corporation

07/ What changes for risk and resilience teams

On one governed topology, the questions that were hardest to answer become standing capabilities.

01/ **Explain concentration** — Show which services carry structural risk from a common dependency, and trace the finding to the paths that create it.

02/ **Challenge redundancy** — Separate genuine alternatives from second paths that share a provider, region, team, control plane or recovery constraint.

03/ **Prioritize treatment** — Compare structural exposures and focus investment on the dependency whose remediation reduces the most consequential exposure.

04/ **Direct scenario testing** — Aim testing at failures with the greatest structural importance rather than building a broad library and hoping.

© 2026 Iroko Technologies. All rights reserved.

05/ **Defend the result**

Retain the evidence, model version, rules and rating lineage needed for challenge, and reopen any past rating exactly as it stood, to show what was true when it was made.

06/ **Keep decisions connected**

Use one operational topology across exposure, impact, simulation and inquiry instead of reconciling separate tools after the fact.

07/ **Validate and calibrate**

Compare assessment results with subsequent incident experience, test whether factors and rating bands distinguish different observed outcomes, and recalibrate the methodology under governance, with structured expert judgment where incident data is thin.

08/ What the measurement does not claim

> It does not replace judgment.

Structural measurement does not replace domain assessment, expert judgment or scenario testing. It adds the exposure created by arrangement: where a component sits, what relies on it, whether alternatives are independent and proven, and how weakness travels through the operating model.

> Reproducibility is not validation.

The engine is deterministic: identical inputs produce identical, decomposable results, and no AI model generates the number. That makes the measurement reproducible, not validated, so the assumptions it rests on, the completeness of the topology, the quality of the assessments and the governance of the method stay visible and open to challenge.

> It is not better than its evidence.

Its evidence can and should improve it: assessment results can be compared with subsequent incident experience to test whether higher ratings distinguish among observed outcomes in incident frequency, severity and impact, and to recalibrate factor weights where the evidence is sufficient. Incident history remains evidence rather than proof, so validation results carry their confidence context.

© 2026 Iroko Technologies. All rights reserved.

09/ Conclusion

This paper began with three questions that mapping alone cannot answer: how exposed an important service really is, where that exposure comes from, and which single failure would hurt it most. Answering them takes a measure the map alone cannot produce: effective exposure, the governed combination of a component's condition and its position in the structure the firm actually runs on.

Neo4j makes that structure computable at enterprise depth. OneScor performs the measurement on it: deep practitioner knowledge of how financial institutions actually operate, formalized into the governed topology, the algorithms that compute the structural score, and the capabilities that turn the score into the answers those questions demand. Together they move the operating model from something a firm can view to something it can interrogate, measure and defend.

> *From documenting dependencies to measuring the exposure they carry.*

In a networked enterprise, structure is part of the risk. Once it is represented, governed and measured, the firm can identify the dependency that matters most, before disruption makes it obvious.

10/ About the collaboration

OneScor, by Iroko Technologies, is a risk and resilience platform built on a governed operational topology. It connects the records that risk and resilience functions maintain one at a time, then computes exposure across them as a standing measure rather than one reconstructed on demand. Every node carries an effective exposure combining its own assessed condition with what it inherits through structure, decomposable back to the paths and factors behind it.

Neo4j is the native graph foundation the operational topology runs on, making traversal, reachability and path analysis direct and scalable at enterprise depth. Together they pair a graph-native foundation with OneScor's governed measurement, methodology and workflow, so institutions can see where exposure concentrates, test the topology against disruption before it happens, and compare interventions by the exposure each one removes. A companion paper, Operational Topology, sets out the underlying argument.

About the authors

> **James Hardy** — Co-founder/CEO, Iroko Technologies

> **Nicholas Hardy** — Co-founder, Iroko Technologies

> **Michael Down** — Global Head of Financial Services, Neo4j

irokotechnologies.com/onescor neo4j.com

© 2026 Iroko Technologies. All rights reserved.

Selected Sources

[1] Basel Committee on Banking Supervision, Principles for Operational Resilience, Principle 4: Mapping interconnections and interdependencies.

[2] Prudential Regulation Authority, SS1/21: Operational resilience, impact tolerances for important business services (2021, updated 2022).

[3] Financial Conduct Authority, PS21/3: Building operational resilience (2021).

[4] Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA), 2022.

[5] Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, Federal Deposit Insurance Corporation, Sound Practices to Strengthen Operational Resilience (2020).

[6] Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, Federal Deposit Insurance Corporation, Interagency Guidance on Third-Party Relationships: Risk Management (2023).

[7] Neo4j Documentation, Graph database concepts: traversals and paths.

© 2026 Iroko Technologies. All rights reserved.