INSIGHT SERIES · GOVERNANCE, RISK & THIRD-PARTY OVERSIGHT

Beyond Your Own Walls

The risk that lives in your supply chain, your vendors, your contractors, and your cloud providers is still your risk. Most organisations are only beginning to reckon with what that means.

Author: E van Doorn, Analysis & Opinion · **Date:** March 2026 · **Topic:** Risk / Compliance

EXECUTIVE SUMMARY

Third-party risk is no longer a peripheral compliance concern — it is one of the most consequential and least-understood exposures facing modern organisations. Incidents at Target, SolarWinds, and a small Florida water utility show how a single vendor, contractor, or software update can become the vector for catastrophic failure. This article examines why visibility, concentration risk, and human factors in vendor relationships demand the same governance rigour organisations apply to their own operations, and sets out what effective third-party oversight actually looks like in practice.

1. Introduction

In January 2021, a water treatment plant in Oldsmar, Florida became, briefly, one of the most alarming stories in American infrastructure security. An intruder had gained remote access to the plant's control systems and attempted to increase the concentration of sodium hydroxide in the water supply to potentially lethal levels. A vigilant operator noticed the cursor moving on his screen and reversed the change within seconds. No one was harmed.

The forensic investigation that followed produced a finding that was, to specialists in operational technology security, entirely unsurprising. The attack had not come through the plant's own systems in any meaningful sense. It had come through a remote desktop application — software installed by a third-party contractor who managed certain functions of the plant's IT environment. The vulnerability was not the plant's. It was the contractor's. The consequence, had the attack succeeded, would have been the plant's entirely.

This is the central paradox of third-party risk in the modern organisation: you can outsource the activity, but you cannot outsource the accountability. The systems, services, data, and processes that flow through your vendors, suppliers, contractors, and technology providers are, from the perspective of your customers, your regulators, and the public, yours. When they fail — when the supplier is compromised, when the contractor makes a catastrophic error, when the cloud provider suffers an outage at precisely the wrong moment — the reputational, financial, and operational consequences land at your door, not theirs.

Most organisations understand this in principle. Far fewer have built the governance and intelligence infrastructure to manage it in practice.

> *"You can outsource the activity, but you cannot outsource the accountability."*

2. The Expanding Perimeter

The scale of the problem is best understood not as a security issue but as a structural one. Over the past three decades, the boundaries of the organisation have dissolved in ways that are economically rational and operationally sensible but that have created a risk surface of extraordinary complexity.

Outsourcing, offshoring, platform dependency, cloud migration, the gig economy, just-in-time supply chains, software-as-a-service — each of these developments transferred activity outside the

organisation while leaving the organisation exposed to the consequences of how that activity is performed. The CFO who outsourced payroll processing to a software vendor in the 1990s was pursuing entirely legitimate efficiency gains. The CISO who discovered, two decades later, that a breach of that vendor's systems had exposed the personal data of every employee the company had ever had was confronting the downstream consequence of a decision that had never been evaluated as a risk decision at all.

The problem has compounded with each wave of technological and commercial change. The modern enterprise does not simply have a list of suppliers. It has a supply chain with multiple tiers, in which its direct suppliers have their own suppliers, who have their own suppliers in turn, each of whom carries a portion of the organisation's risk exposure without any direct relationship, contract, or visibility. It has technology providers whose systems are so deeply embedded in its own operations that disentanglement is no longer practically possible. It has contractors and professional services firms who move in and out of its most sensitive environments with credentials and access levels that frequently exceed what their function requires. And it has, in many cases, very little systematic visibility into any of it.

2.1 The Regulatory Response

The regulatory community has been slow to catch up, but it is catching up. Three developments illustrate the direction of travel:

1. The European Union's Digital Operational Resilience Act, which came into force for financial services firms in January 2025, imposes explicit and detailed requirements on the oversight of critical third-party technology providers — including, in certain cases, direct regulatory oversight of the providers themselves.
2. The SEC's cybersecurity disclosure rules, adopted in the United States in 2023, require listed companies to disclose material cybersecurity incidents, including those that arise through third parties.
3. The UK's operational resilience framework, developed by the Bank of England and the Financial Conduct Authority, places the management of outsourcing and third-party dependencies at the centre of firms' resilience obligations.

> **Key Point:** The direction of travel is clear. The expectation — from regulators, from boards, and increasingly from customers and counterparties — is that organisations know who they depend on, understand what risks those dependencies carry, and can demonstrate that they are managing them with something more rigorous than periodic questionnaire and annual contract review.

3. The Visibility Problem

The first and most fundamental challenge in third-party oversight is simply knowing what you have. This sounds trivially simple. It is not.

Ask the chief procurement officer of a large organisation how many active third-party relationships their organisation maintains. The number they give you will typically reflect what is visible through the formal procurement process — the suppliers who have been properly onboarded, the contracts that have been reviewed by legal, the vendors who appear in the accounts payable system. It will not reflect the shadow IT that operating units have acquired independently, the software-as-a-service tools that teams have subscribed to on a corporate credit card, the contractors engaged directly by individual departments, or the subcontractors and fourth parties that the organisation's tier-one suppliers have engaged on their behalf.

The gap between the formal vendor register and the actual third-party footprint is, in most large organisations, significant. Research conducted across multiple sectors consistently finds that organisations undercount their third-party relationships by a substantial margin — in some cases, the true number is two or three times what the formal register suggests.

3.1 Risk-Proportionate Tiering

This is not, in the main, the result of negligence or bad governance. It is the result of the way modern organisations actually work: distributed, decentralised, responsive to operational need, and resistant to the friction that comprehensive procurement governance imposes. The business unit that needs a software tool to solve an immediate problem does not want to wait three months for it to complete the vendor onboarding process. The project manager who needs a specialist contractor for two weeks does not see why a fifteen-page third-party risk questionnaire is proportionate to the engagement. These are not unreasonable positions. They reflect the genuine tension between the speed at which business operates and the thoroughness that risk management requires.

Managing that tension is one of the central practical challenges of third-party oversight. The organisations that manage it best do not attempt to impose the same level of scrutiny on every third-party relationship. They invest, instead, in developing a rigorous, risk-proportionate tiering system — a framework that distinguishes between the cloud provider on whose systems the entire business depends and the stationery supplier whose relationship with the organisation involves essentially no risk transfer. The scrutiny applied to the former should be intensive, continuous, and senior. The scrutiny applied to the latter should be minimal, automated where possible, and largely administrative.

The organisations that manage this tension best calibrate scrutiny explicitly by tier:

- Intensive, continuous, and senior scrutiny for the providers on whom the entire business depends.
- Minimal, largely automated, and administrative oversight for relationships that involve essentially no risk transfer.

> **Key Point:** The critical insight is that visibility does not mean uniformity. It means knowing enough about every relationship to know how much attention it deserves.

4. Concentration Risk: The Dependency You Did Not Notice Building

Among the risks that third-party exposure creates, concentration risk is perhaps the most structurally underappreciated. Concentration risk arises when an organisation's dependence on a small number of third parties — or, more insidiously, when the market's dependence on a small number of third parties — creates a single point of failure of systemic significance.

The cloud infrastructure market provides the clearest contemporary illustration. Three providers — Amazon Web Services, Microsoft Azure, and Google Cloud — account for the substantial majority of cloud services consumed globally. Organisations across every sector have migrated their operations onto infrastructure provided by one or more of these providers in pursuit of the efficiency, scalability, and capability that cloud platforms offer. The benefits are real. The concentration risk is equally real and considerably less discussed.

When a major cloud provider experiences a significant outage, the consequences are not confined to the provider's own customers. They propagate through the supply chains, the service providers, and the technology stacks of an enormous and often invisible web of downstream dependencies. The organisations that are affected may not even know that their operations depend on the affected provider, because their dependence is mediated through a software vendor or a payment processor or an analytics platform that itself runs on cloud infrastructure that has gone dark.

This is the fourth-party problem — the risk that lives not in your direct suppliers but in your suppliers' suppliers — and it is among the least well-managed dimensions of third-party risk in most organisations. The standard third-party risk management framework addresses tier-one relationships with some rigour. It typically addresses tier-two relationships with considerably less,

and tier-three relationships with essentially none, on the reasonable grounds that pursuing visibility that deep is practically difficult and resource-intensive.

The weakness of this reasoning became apparent in July 2024, when a defective software update pushed by the cybersecurity firm CrowdStrike caused an estimated 8.5 million Windows devices globally to display the Blue Screen of Death simultaneously. The organisations affected had not, in the main, contracted directly with CrowdStrike. The software was embedded in their IT environments through the systems of managed service providers, cloud platforms, and software vendors who themselves used CrowdStrike as part of their security stack. The dependency was invisible — until it was not.

5. Case Study I: The Supplier Who Brought Down a Retailer — Target, 2013

In the winter of 2013, Target Corporation suffered a data breach that exposed the payment card details of approximately 40 million customers and the personal information of a further 70 million. It was, at the time, among the largest retail data breaches in history. The reputational and financial cost was enormous — the company's chief information officer and chief executive both ultimately resigned, and the total cost of the breach has been estimated at over $200 million, exclusive of longer-term brand damage.

The breach did not originate inside Target. It originated with a third-party contractor — a refrigeration and HVAC firm based in Pennsylvania, called Fazio Mechanical Services, which had been engaged to provide facilities management services to Target stores across the United States. Fazio had been granted network credentials to access Target's systems — credentials that allowed it to submit electronic invoices, manage contracts, and monitor energy consumption and temperature in Target's facilities. The credentials were legitimate. The access was, in retrospect, excessive.

Attackers who had compromised Fazio's systems used those credentials to gain an initial foothold in Target's environment. From there, they moved laterally through Target's network — a process made possible by the absence of network segmentation that would have prevented a vendor with facilities management credentials from accessing systems relevant to payment processing. They eventually installed malware on the point-of-sale systems in Target's stores, capturing card data as it was swiped by customers.

The post-mortem revealed multiple failures in Target's third-party oversight framework. The risk associated with Fazio's network access had not been assessed in a manner proportionate to the sensitivity of the systems that access could potentially reach. The access rights themselves had not been reviewed or restricted to what the relationship actually required. The network architecture did

not impose the segmentation that would have limited the blast radius of a third-party compromise. And the monitoring systems that might have detected the lateral movement once it began were not configured to generate alerts that reached the people with authority to act on them.

What made the Target breach instructive was not its technical complexity — the attack was not, by the standards of sophisticated adversaries, particularly advanced. What made it instructive was the way it illustrated the systematic gap between the formal governance of third-party relationships and the actual risk that those relationships carry. The vendor had been engaged through proper channels. The contract was in place. The relationship existed, in the organisation's records, as a managed and understood dependency. The risk it carried was invisible to everyone who might have managed it.

6. Case Study II: When the Ecosystem Fails — SolarWinds, 2020

The SolarWinds compromise, disclosed in December 2020, is the canonical illustration of what happens when third-party risk operates at ecosystem scale. The attack was, in its mechanics, a software supply chain compromise: adversaries — subsequently attributed to the Russian foreign intelligence service, the SVR — had inserted malicious code into the build process for SolarWinds' Orion software, a widely used IT performance monitoring platform. Updates to Orion pushed to customers between March and June 2020 contained the malicious code, known as SUNBURST.

The scale of the exposure was extraordinary. SolarWinds had approximately 18,000 customers who downloaded the compromised update. Those customers included the United States Treasury, the Department of Homeland Security, the Department of State, the National Institutes of Health, and a significant number of Fortune 500 companies. The attackers had, in effect, used a trusted third-party software provider as an unwitting distribution mechanism for a tool that gave them persistent, privileged access to the internal networks of organisations that had done nothing wrong in their own environments.

This is the defining characteristic of a supply chain attack: the victim's own defences are irrelevant. The malicious code arrived signed and verified, embedded in an update from a vendor that customers had every reason to trust. No amount of conventional security hygiene at the customer level would have prevented the initial infection. The vulnerability was in the supply chain, not in the customer's own walls.

The implications for third-party oversight are profound. They suggest that the traditional model of third-party risk management — which focuses primarily on assessing whether a vendor maintains adequate security controls in its own environment — is necessary but insufficient. It addresses the risk that your vendor is compromised because of its own weaknesses. It does not, in itself, address

the risk that your vendor becomes, through no particular failing of its own, a vector for an adversary who targets it precisely because of the access its customers have granted it.

Managing this class of risk requires a different kind of thinking: about the integrity of the software and services entering the organisation's environment, about the monitoring that would detect anomalous behaviour from trusted sources, and about the fundamental question of how much access any third party — however trusted — should ever hold. The principle of least privilege — granting access to only what a third party needs, and revoking it when the need has passed — is well understood in theory. The SolarWinds case demonstrated, at considerable cost, how far from routine implementation it remains in practice.

7. What Good Looks Like

Effective third-party oversight is not a single programme or a single function. It is the integration of several capabilities that, in most organisations, currently exist in fragments — procurement, legal, cybersecurity, operational resilience, and compliance each managing a portion of the third-party landscape with limited coordination and significant gaps.

The organisations that lead in this space — and they are identifiable, concentrated in financial services, defence, and sectors with long regulatory experience of outsourcing risk — share several characteristics.

In practice, this means:

- They have genuine inventory. Not just a vendor register, but a living map of third-party relationships that includes tier-two and tier-three dependencies where the risk is material, that is continuously updated as relationships are established and terminated, and that is connected to the organisation's operational and data architecture well enough to understand what each relationship actually touches.
- They apply risk-proportionate scrutiny. The depth and frequency of oversight is calibrated to the criticality and sensitivity of each relationship — intensive and continuous for the providers on whom the organisation is operationally dependent, lighter and more automated for the relationships that carry limited risk. The tiering is explicit, documented, and consistently applied.
- They manage concentration actively. They know which third parties — and which third-party ecosystems — represent single points of failure, and they have thought carefully about what mitigation is possible. In some cases that means maintaining fallback arrangements. In

others it means accepting the concentration risk but holding it at board level as a named, owned, and monitored exposure.

- They treat contracts as the beginning of oversight, not the end. The contract with a critical supplier is not the primary tool for managing the relationship's risk. It is the framework within which continuous operational engagement, performance monitoring, and periodic deep-dive assessment take place. The organisations that discover they have a problem only when they try to invoke a contract clause have already lost the benefit of third-party oversight.
- And they have thought about exit. The capacity to move away from a third-party relationship — to switch providers, to insource a capability, to wind down a dependency — is one of the most neglected dimensions of third-party oversight. It is also one of the most practically important. An organisation that cannot credibly exit a relationship has implicitly surrendered a substantial portion of its leverage over the provider and a substantial portion of its resilience in the event that the relationship fails.

8. The Human Element

There is a dimension of third-party risk that tends to receive less attention than the technological and contractual dimensions, and that may in practice be more consequential: the human element.

Contractors, consultants, and professional services staff who work within an organisation's environment are, functionally, part of that organisation during the period of their engagement. They have access to systems, to data, to conversations and decisions. They often know more about the organisation's vulnerabilities than the organisation's own staff — precisely because they are brought in to work on the most complex and sensitive problems. And they operate, in most cases, with a level of oversight that reflects neither the sensitivity of their access nor the risk that their engagement creates.

The insider risk that contractors carry is not primarily a function of malicious intent — though that cannot be entirely excluded. It is primarily a function of the structural fact that their loyalty, their incentives, and their employer obligations all lie outside the organisation they are serving. The data they encounter during an engagement lives in their memory, and often in their files, after the engagement ends. The vulnerabilities they have identified in the course of their work may or may not be known to the organisation after they leave.

Managing this risk requires both structural and relational interventions. Structurally: appropriate access controls, clear data handling requirements, offboarding processes that recover credentials and assets, and monitoring that does not distinguish between employees and contractors in ways that create blind spots. Relationally: the cultivation of genuine professional relationships with key third-

party staff that create the mutual understanding and trust in which concerns are raised rather than concealed, and in which the standards of behaviour expected of the organisation's own people are extended naturally to those working alongside them.

9. The Regulatory Horizon

Regulation of third-party risk is moving, globally, in one clear direction: toward greater specificity, greater accountability, and greater consequence for organisations that cannot demonstrate meaningful oversight of their critical dependencies.

The EU's Digital Operational Resilience Act represents the most comprehensive legislative expression of this trend in financial services. It requires firms not only to manage their third-party IT relationships but to register critical providers with regulators, to conduct thorough due diligence and periodic testing of those relationships, and to ensure that contracts with critical providers contain minimum substantive provisions. It also, significantly, creates a framework for direct regulatory oversight of critical third-party providers themselves — a recognition that the risk posed by concentrated dependencies in the financial system cannot be fully addressed by regulating only the firms that use them.

Similar logic is beginning to appear in healthcare, in defence contracting, in critical infrastructure regulation, and in the emerging frameworks around AI system governance, where the chain of dependencies — foundation model providers, fine-tuning services, deployment platforms, integration contractors — creates third-party risk of a kind that the existing frameworks were not designed to address.

The organisations that are building rigorous third-party oversight now are not merely responding to current regulatory requirements. They are positioning themselves ahead of a regulatory evolution that is, in its broad direction, entirely predictable. The specific requirements will vary by jurisdiction and sector. The underlying expectation — that organisations can see, understand, and manage the risk that lives beyond their own walls — will not.

Conclusion

The era in which an organisation's risk perimeter could be meaningfully defined by what it owned and directly controlled ended some time ago. The replacement era — in which the perimeter is defined instead by what the organisation depends on, and by what can harm it through those dependencies — is now the operating reality for virtually every institution of any scale.

The inescapable conclusion of that reality is that third-party oversight cannot remain a compliance function, a procurement sub-process, or an annual questionnaire exercise. It must be a genuine organisational capability — resourced, governed, continuously exercised, and integrated with the risk intelligence and governance frameworks that determine how the organisation sees and manages risk as a whole.

The risk that lives in your supply chain, your vendors, your contractors, and your cloud providers is still your risk. The organisations that have grasped this — not as a regulatory obligation but as a strategic reality — are the ones building the visibility, the relationships, and the resilience to manage it. The ones that have not are, with a reliability that the post-mortem record makes difficult to dispute, the ones that will be explaining, at some point, how something that was not technically their fault turned out to be entirely their problem.

SOURCES

This article draws on the Target Corporation post-breach analysis published by the US Senate Commerce Committee (2014), the Cybersecurity and Infrastructure Security Agency's advisory on the SolarWinds compromise (2020–2021), the EU Digital Operational Resilience Act (Regulation 2022/2554), and the NATO Cooperative Cyber Defence Centre of Excellence's documentation of the Oldsmar water treatment incident. The CrowdStrike outage figures are drawn from Microsoft's preliminary assessment published in July 2024. Information on RailGuard's platform capabilities is drawn from published documentation at assessum.com.