Third-party risk, critical supply chains and cyber resilience

Summary
Third-party risk has become a defining resilience challenge for modern organisations. Few enterprises operate within clearly bounded environments anymore. Critical business processes, IT services, data flows and customer-facing capabilities increasingly depend on a broad ecosystem of software providers, cloud platforms, managed service providers, professional services firms and AI vendors. Fourth-party dependencies are hidden further down the supply chain. As a result, organisations are no longer only as secure as their own controls, but as secure as the wider ecosystem on which they depend.

This session explored how that reality is reshaping third-party cyber risk management. The central message was clear: traditional third-party risk management approaches, built around periodic assessments and static assurance, are no longer sufficient in an environment where threats, dependencies, software components and exposures change continuously and quickly. The discussion therefore moved beyond compliance-oriented vendor assurance toward a more resilience-based model, focused on visibility, validation, prioritisation and adaptive response.

The new reality: organisations depend on ecosystems
Biswajit Behera opened the session by positioning third-party risk as an ecosystem problem. Organisations now rely on interconnected third parties across SaaS, cloud, managed services, AI solutions, and professional services, while also inheriting risks from fourth parties, open source libraries, software supply chains and hardware-software dependencies. This means that the effective attack surface extends far beyond organisational boundaries.

In that environment, cyber resilience cannot be understood only in terms of internal maturity. Even where internal controls are relatively strong, weaknesses in suppliers, service providers, or software dependencies may still create direct operational and security consequences. Third-party risk is therefore no longer peripheral. It has become part of core business resilience, because disruption at a key supplier may affect operations, data integrity, customer trust, compliance and continuity simultaneously.

This broader framing aligns with the wider market view that third-party risk management is being reshaped by AI adoption, outsourced digital services, and increasingly opaque vendor ecosystems. Traditional TPRM methods were not designed for questions such as embedded AI use, model governance, data lineage or continuously changing software dependencies. Biswajit challenged the audience with five strategic questions.

> An organization is only as secure as the ecosystem it relies upon.
>
> 
>
> **Biswajit Behera**
> ING

Five strategic questions for organisations
One of the strongest introductions for discussion in the session was a set of five questions presented to provoke the attendees reflection. Together they help shift the conversation from process completion to resilience readiness:

1. **How much confidence should we place in external assurance?** External reports and certifications remain useful, but they are not the same as real-time trustworthiness. Organisations need to understand both the value and the limits of attestation-based assurance.
2. **Do we understand our critical third- and fourth-party dependencies?** Many organisations know their direct suppliers only partially and have even less visibility into deeper dependencies. Yet this is where concentration risk and hidden systemic fragility often sit.
3. **Could we detect supplier compromise before disclosure?** If the answer is no, then the organisation depends largely on the supplier's own detection and transparency timelines, which may be too slow in a fast-moving incident.
4. **Are we measuring supplier compliance or supplier resilience?** These are not the same. A supplier may appear compliant on paper while still being operationally fragile or slow to detect and contain threats.
5. **What would fail if a critical supplier were compromised tomorrow?** This question forces a business continuity and resilience lens. It shifts attention from abstract control review to practical impact on services, operations, revenue, and recovery.

These questions are valuable because they move TPRM away from a documentation exercise and toward strategic dependency management. They also underline that resilience

is not only about preventing all incidents, but about understanding which failures matter most and preparing to absorb, contain, and recover from them.

Attackers target trusted relationships

A key theme throughout the session was that attackers are increasingly exploiting trusted relationships rather than attempting direct intrusion through hardened organisational perimeters. Software dependencies, shared tooling, valid accounts, federated access, and supplier connections have become attractive entry points because they offer scale, legitimacy, and often weaker visibility.

The ING presentation captured this clearly by mapping supply chain compromise, trusted relationship abuse, valid account abuse, ransomware, and data exfiltration as relevant attack patterns. The examples discussed showed that this is not a theoretical concern. Supply chain attacks and vendor-linked incidents can spread quickly across many customers, repositories, users, or business environments, making the impact systemic rather than isolated.

This reflects a wider trend in the market, where growing dependence on third parties is intensifying both the number of relationships to assess and the complexity of the risks those relationships introduce. As dependency expands, so does the challenge of identifying what is truly critical, where hidden dependencies sit, and how quickly compromise could propagate through the ecosystem.

> Identifying vulnerabilities at scale generates noise; prioritizing them generates security.
>
> 
>
> **Stefan van Vlerken**
> Okta

The assurance gap in traditional TPRM

One of the clearest conclusions from the session was that there is now an assurance gap between traditional TPRM practices and the real dynamics of third-party cyber risk. Conventional approaches still rely heavily on questionnaires, external assurance reports such as SOC 2, static risk ratings, and annual or periodic reviews. While these instruments remain useful, they offer only point-in-time confidence.

That is increasingly insufficient. They may confirm that a vendor had certain controls, documents, or attestations at a particular moment, but they often fail to reveal active compromise, emerging vulnerabilities, credential theft, supply chain attacks, or rapidly changing exposure. In other words, they tend to measure declared compliance more effectively than operational resilience.

This gap matters because the risk environment now changes much faster than review cycles. Threat actors adapt continuously, software changes frequently, external attack surfaces evolve, and AI is accelerating both the pace and scale of abuse. A yearly review process cannot keep up with a threat landscape that changes weekly or daily.

The same limitation is recognised in broader TPRM and AI governance discussions. Updated third-party oversight requires more than checkbox diligence; it requires organisations to integrate AI-oriented and threat-oriented controls into existing risk frameworks and to shift TPRM from a reactive gatekeeper into a proactive enabler of secure adoption.

Trust must be validated continuously

The session therefore argued for a shift from static assurance to continuous validation. Trust in third parties should not be assumed once a contract is signed or a questionnaire is completed. It needs to be tested and refreshed continuously through dynamic signals that reflect real operational conditions.

The ING material described this as continuous monitoring supported by vendor risk intelligence derived from multiple categories of indicators. These include threat intelligence signals such as credential leaks, ransomware targeting, and dark web mentions; software and dependency signals such as critical vulnerabilities, dependency changes, and end-of-life components; identity-related signals such as MFA abuse, privilege escalation, or federated access anomalies; and external exposure indicators such as DNS changes, certificate changes, known exploited vulnerabilities, and attack surface observations.

The value of this approach is that it provides a more current and evidence-based view of vendor risk. Rather than relying only on vendor declarations, organisations can observe changes in posture, indicators of compromise, or patterns of weakness as they emerge. This makes TPRM more adaptive and better aligned with the speed of the threat environment. This direction is also increasingly reflected in vendor and industry discussions around AI-enabled TPRM. Continuous monitoring, dynamic scoring, automated document analysis, and real-time external signals are presented as ways to reduce manual burden while improving visibility and timeliness. Those capabilities should still be treated critically, but the underlying shift from static review to continuous monitoring is well founded.

More effective prioritisation: from exposure to control
Another important contribution in the session came from the presentation from Stefan van Vlerken on moving from exposure to control. He stated that organisations often generate much noise and spend much effort on findings that appear severe in theory but are not actually the most urgent in practice. Discovery alone produces scale, but not necessarily clarity. Validation without broad discovery, on the other hand, leaves blind spots. The combination of the two is needed to understand what is truly at risk.

Stefan presented a four-step model: discover, validate, prioritise and remediate. This sequence matters because it prevents organisations from treating all findings equally or allowing raw severity scores to dictate action without operational context. A vulnerability with a high CVSS score but no exploit activity, no weaponisation, no malware use and no meaningful exposure may be less urgent than a technically lower-scored issue that is actively exploited, trending, or tied to critical business assets.

The presentation showed how threat context materially improves prioritisation. Signals such as weaponised exploits, inclusion in known exploited vulnerability catalogues, malware use, threat actor activity, and business criticality make it possible to focus on the relatively small subset of exposures that genuinely matter to the organisation. This creates substantial efficiency gains while also improving risk outcomes, because scarce remediation capacity is directed where it has the highest resilience impact.

The shift from risk management to resilience
The overall conclusion of the session was that third-party cyber risk management is evolving from a primarily compliance- and risk-oriented function into a resilience-oriented discipline. In the earlier model, the emphasis was often on reactive questionnaires, periodic assurance, and administrative control checks. In the emerging model, the emphasis is on managed risk tiering, integrated continuous monitoring, threat-informed decision making, and cyber supply chain governance.

This is not a rejection of governance or compliance. Rather, it is a recognition that governance must become more adaptive and better informed by operational reality. A resilient approach combines structured due diligence with continuous telemetry, threat intelligence, software dependency visibility, and response planning. It also requires closer alignment between procurement, cyber security, risk, legal, architecture and business continuity functions.

The session's message was therefore both practical and strategic. Practical, because it identified concrete improvements such as continuous monitoring, dependency mapping, threat-informed prioritisation, and dynamic validation of trust. Strategic, because it argued that third-party resilience is now inseparable from organisational resilience itself.

Concluding reflection
The discussion made clear that third-party risk can no longer be managed effectively through static assurance alone. The combination of digital interdependence, software supply chain complexity, AI-enabled services, and increasingly sophisticated adversaries has created a threat environment in which trust must be earned repeatedly, not assumed once.

For highly resilient organisations, the task is therefore not simply to manage vendor risk as a compliance process, but to build a sustained capability to understand dependencies, detect change, prioritise what matters, and respond before disruption cascades across the ecosystem. In that sense, the future of TPRM is not only about better oversight. It is about building operational resilience in a world where the boundaries of the organisation are no longer the boundaries of its risk.

You should broaden testing scenario's to your whole ecosystem and find hidden vendor risks before attackers do.

Suzanne Janse
Program Director HRO