The Missing Measure in Third-Party Information Risk

Making third-party information risk governable, comparable, and transferable

Daniel Nutkis
Founder and Executive Chairman, HITRUST
May 2026

EXECUTIVE SUMMARY

Third-party information risk is now an enterprise management issue. Vendor, cloud, platform, processor, and subcontractor dependence creates exposure beyond cyber failure, including operational disruption, regulatory risk, contractual loss, revenue impact, reputational harm, uninsured financial loss, and resilience gaps.

Current tools show signals, not the full risk picture. Cyber scores, monitoring tools, questionnaires, certifications, audits, contracts, and insurance evidence all help, but they vary in scope, rigor, timing, and assumptions. A completed review does not necessarily produce a clear measure of residual exposure.

The missing measure is a trusted way to convert fragmented evidence into comparable residual risk. An assured, standardized methodology or index that can normalize third-party evidence and business context into decision-ready insight, is needed.

Better measurement changes governance. With a common risk language, leaders can set thresholds, approve or reject vendors, manage exceptions, identify portfolio concentrations, benchmark posture, evaluate financial impact, and determine what risk is retained or transferred.

Board-level takeaway: third-party risk needs the same measurement discipline as other enterprise risks. A common number is useful only if the methodology behind it is documented, governed, validated, and trusted enough to support management decisions, executive oversight, and market reliance.

Third-party information risk has become an enterprise management problem.

Enterprises rely on vendors, suppliers, platforms, processors, cloud providers, subcontractors, and other service organizations to run critical operations, handle sensitive and regulated information, support customer outcomes, and maintain business continuity.

That dependence creates third-party information risk, which is the possibility that information outside the enterprise's direct control is not protected, governed, processed, shared, used, or recoverable in a way that aligns with the organization's risk appetite, legal obligations, contractual commitments, and continuity expectations.

Cybersecurity failure is only one expression of this exposure. Third-party information risk can also create operational disruption, privacy impact, regulatory exposure, contractual loss, revenue impairment, reputational harm, uninsured financial loss, and reduced resilience.

The central question is whether the organization can understand the residual exposure created by third-party dependence and use that understanding to make better decisions, in terms of which vendors to approve, which risks to remediate, which exceptions to accept, which exposures to aggregate, which risks to transfer, and how performance compares across vendors and peers.

Most large enterprises have built Third-Party Risk Management programs to respond to this challenge. In principle, those programs should identify the full vendor population, tier vendors by inherent risk, and evaluate each relationship based on data sensitivity, business criticality, connectivity, regulatory exposure, geography, substitutability, and operational dependency. They request and review questionnaires, certifications, audit reports, and submitted evidence, then evaluate control gaps, require contractual commitments and insurance, and route exceptions through approval workflows.

These activities are necessary, but they are not sufficient.

In practice, many organizations do not have the time, staffing, or process capacity to evaluate the full vendor population adequately. Teams often focus on a subset of vendors, the most visible, highest risk, most business-critical, or most urgent relationships. That leaves unknown exposure across the broader ecosystem and makes it difficult to know whether the vendor population has been evaluated consistently.

Current cyber scores are useful, but they are not the missing measure.

External cyber ratings, attack-surface monitoring, threat-intelligence indicators, and monitoring tools can help identify exposed assets, hygiene issues, breach signals, and changes in posture at scale. They are valuable inputs for triage, prioritization, and ongoing monitoring.

But third-party information risk requires more than an outside-in view of cyber posture. The enterprise needs to know whether the specific service, data flow, control scope, contractual commitments, insurance risk transfer, business criticality, and assurance evidence support a decision within risk appetite. A cyber score may identify signals of concern, but it does not by itself distinguish control effectiveness, assurance confidence, residual risk, retained exposure, transferred exposure, and plausible financial impact.

The missing measure is an assured method for converting all relevant inputs, be it outside-in signals, certifications, audits, questionnaires, contracts, insurance, remediation status, business context, and dependency data into a comparable measure of residual and retained exposure.

Questionnaires, certifications, audit reports, remediation updates, control scorecards, and external signals can all be useful. They also differ in scope, rigor, timing, assumptions, exclusions, limitations, and format.

Organizations therefore need a standardized way to consume, normalize, interpret, and verify those inputs. Without that foundation, a completed review may confirm process completion while leaving the enterprise without an accountable measure of residual exposure.

That operational gap becomes a governance problem when organizations define and apply risk thresholds. Effective governance requires the enterprise to determine what level of third-party information risk is acceptable for a given vendor relationship and apply that standard consistently across vendors, reviewers, and the broader ecosystem.

Those thresholds cannot be generic. They should reflect inherent risk, data sensitivity, business criticality, regulatory exposure, geography, substitutability, contractual protections, compensating controls, and residual risk after treatment. A vendor supporting a critical business process or handling regulated information warrants a different tolerance than a low-risk supplier with limited access or dependency.

This requires more than evidence collection. It requires a consistent way to measure residual information risk and compare that measure to the organization's risk appetite. Without a stable unit of measure, threshold decisions drift toward reviewer experience, business urgency, negotiation leverage, available documentation, or local interpretation.

Exception decisions require the same discipline. When a vendor falls outside normal thresholds, the enterprise must compare the risk of proceeding with the business impact of delay, rejection, replacement, or redesign. Decision-makers need to understand what residual exposure remains after controls, assurance evidence, contractual protections, insurance, and remediation commitments are considered. They also need to understand the business impact on customer obligations, revenue opportunities, regulatory requirements, transformation initiatives, continuity needs, and available alternatives.

Business urgency may justify proceeding outside normal tolerance. But the accepted exposure should be explicit. An exception should identify the nature and driver of the tolerance gap, the business consequence of delay or rejection, the alternatives available, the residual exposure accepted, the expected duration, the accountable business owner, and the mitigation or transfer plan.

Individual decisions can also accumulate into portfolio risk. One exception may be manageable when the exposure is understood and the mitigation plan is credible. Many similar decisions across vendors, data types, geographies, business functions, technologies, subcontractors, or control domains can expand the attack surface and create material concentration risk.

Most TPRM programs are built to move individual vendors through review workflows. Fewer are built to quantify cumulative residual risk, estimate the financial impact of accepted exceptions, or identify concentrations across the vendor portfolio and ecosystem. As more enterprise information risk sits outside direct enterprise control, organizations need a portfolio view that is comparable, aggregable, and decision-ready.

That portfolio view should show residual risk by vendor, business process, data type, geography, control domain, critical service, subcontractor dependency, technology dependency, and business unit. It should also show how exposure changes when controls are remediated, certifications expire, incidents occur, insurance limits change, or the enterprise becomes more dependent on a provider.

Benchmarking also depends on a common measurement language. It helps organizations understand whether their third-party information risk posture falls within expected ranges, whether thresholds are calibrated appropriately, whether exception levels are unusual, and whether remediation priorities align with the exposures that matter most.

Today, internal scores use different scales, assessment rigor varies by team and vendor tier, and the same vendor profile may be interpreted differently across organizations. A "high-risk" designation in one organization can mean something very different in another.

Risk transfer is another reason measurement matters. It can allow an organization to proceed with a vendor that would otherwise exceed normal tolerance, support a strategically important relationship, or reduce the financial exposure the enterprise must retain. But transfer changes who is expected to bear some portion of the financial consequence. It does not eliminate the underlying operational or information risk.

Current models can create false comfort. A vendor's cyber insurance policy may leave one client's losses inadequately addressed. Contractual indemnity may be limited by liability caps, exclusions, enforceability, dispute process, vendor solvency, and the funding available behind the promise. Certificates and contract terms are useful inputs. Residual-risk measurement still requires separate evaluation.

The missing layer is a standardized, assured, and accountable method for converting inconsistent third-party evidence into a comparable measure of residual information risk. Expressed through an index, that methodology would provide a common reference scale for vendor decisions, threshold governance, exception management, portfolio aggregation, peer benchmarking, financial impact analysis, and risk transfer.

A common number without a trusted methodology would become another subjective input.

To be credible, the methodology behind the index must be documented, governed, validated, and applied consistently. An effective measurement foundation should include five elements:

- **Common risk language** to distinguish control effectiveness, assurance confidence, residual risk, retained exposure, transferred exposure, and plausible financial impact.
- **Evidence normalization and assurance weighting** to define how questionnaires, certifications, audits, contracts, insurance evidence, remediation updates, external cyber signals, and attestations are mapped and interpreted.
- **Control-to-risk and residual-risk methodology** to connect evidence to the risks it reduces and adjust inherent exposure for verified controls, evidence gaps, compensating controls, unresolved exceptions, and business context.
- **Decision and portfolio outputs** to support approval, escalation, documented acceptance, rejection, aggregation, benchmarking, financial impact analysis, and risk transfer.
- **Governance, validation, and reliance** to ensure the methodology is versioned, calibrated, reviewed, quality-controlled, and supported by a credible standards and assurance model.

A standardized and universal index should help organizations see what residual exposure remains after controls, contracts, remediation, and insurance are considered. It should help reviewers interpret vendor artifacts consistently, distinguish retained and transferred exposure, identify portfolio concentrations, support benchmarking and calibration, and move quickly when risk is understood and acceptable while escalating evidence gaps, control weaknesses, and material residual exposure.

The next stage of third-party information risk management requires a universal, assured measurement methodology that turns fragmented third-party evidence into comparable residual-risk insight. Such an index would support better vendor decisions, threshold governance, exception management, portfolio aggregation, benchmarking, financial impact analysis, and risk transfer.

Consistent measurement makes third-party information risk governable, comparable, scalable, benchmarkable, and transferable. The index matters because it helps organizations solve that problem with a methodology trusted enough to support management decisions, executive oversight, and market reliance.

IMPORTANCE OF THE METHODOLOGY

A common number, by itself, is not enough. The credibility comes from the method behind it. The methodology needs:

- **A common risk language**
- **Normalized evidence and assurance weighting**
- **A control-to-risk and residual-risk model**
- **Decision and portfolio outputs**
- **Governance and validation**