The hidden weakness in third-party cyber risk transfer

Large enterprises invest heavily in Third-Party Risk Management (TPRM) programs: from tiering vendors based on risk assessments to requiring cyber liability insurance, collecting certificates of insurance, and verifying limits. Despite these efforts, structural weaknesses in both the traditional cyber insurance model and TPRM remain largely unaddressed. This creates a material and previously under-appreciated exposure for the enterprise itself.

The issue is not that enterprises ignore third-party risk. It is that a vendor's cyber insurance policy, the market's default risk transfer mechanism, isn't designed to protect enterprise clients in the way many organizations assume. A vendor's policy is critical and will likely help the vendor respond to a breach appropriately, pay legal their costs, manage notifications and restore operations. But it does not guarantee sufficient aggregate limits for customer-specific protection, especially when a vendor incident creates losses across many clients simultaneously.

Enterprises rely on vendors for critical systems, data processing, infrastructure, security services, payment flows, and operational continuity. One vendor incident can affect dozens, hundreds, or thousands of downstream customers. Yet the insurance structures used to manage that unique cyber exposure for enterprise clients have not sufficiently evolved from their MPL and D&O policy form roots. The quintessential issue into today's cloud-centric business environment is interdependency; yet the insurance products used to manage that exposure today are very likely insufficient, and worse, create a false sense of financial security.

Three gaps stand out: shared and inadequate limits, inconsistent policy language, and recovery that often depends on litigation.

Cyber insurance certificates obfuscate the shared-limits reality

Most enterprise client contracts require vendors to "prove" they maintain a specific aggregate limit of cyber liability insurance, the value usually determined by the nature and/or criticality of a vendor's services. The requirement is met once a vendor produces their certificate of insurance, demonstrating the specified aggregate limit of liability (often between $1 million and $10 million). But that number is almost-universally viewed in isolation. The cert process simply ignores the question that is perhaps most important in determining the adequacy of vendor cyber limits: how many other customers are relying on the same aggregate limit of insurance stated on the cert?

A vendor may support 50 or more enterprise clients under just one aggregate cyber limit. That limit is not dedicated to any individual customer, but each believes they have access to the limit displayed on the certificate. For a contained incident, that aggregate may suffice. For a ransomware event, supply-chain compromise, outage, or breach affecting many customers at once, it becomes fragile. Early claimants can erode the available limit quickly, leaving others with little or no practical recovery. If the vendor's policy is exhausted, the enterprise may have to absorb the loss or turn to its own cyber insurance. That may solve the immediate issue, but it can affect future renewals through higher premiums, larger retentions, narrower coverage, or more restrictive terms.

This is not dependable risk transfer. It is essentially a shared risk pool in which vendor clients are unknowingly swimming.

This dynamic stands in stark contrast to another well-established risk-transfer mechanism: surety bonds used in commercial construction. In construction, the contractor pays the premium, but the project owner receives direct protection. No project owner would ever accept a surety bond issued for, and shared with, an entirely unrelated project; each bond is project-specific and ring-fenced. Enterprise clients should expect the same clarity in vendor cyber risk transfer. The implicit question arises: If you would never allow a contractor to leverage a shared surety bond, why would you accept a vendor's shared cyber policy?

A major 2023 supply-chain cyber incident shows the problem clearly: thousands of organizations and tens of millions of individuals were affected, while the vendor reportedly carried less than $20M in cyber limits. The broader economic impact was estimated in the billions when factoring per-record breach costs, notification, remediation, and litigation. The lesson is structural: one shared policy cannot reliably protect every affected enterprise in a systemic, multi-client loss.

Certs aren't certainty

The second gap is the lack of policy standardization in cyber.

Cyber insurance does not have the same level of standardization as many property or general liability products. Carriers use different policy language, definitions, triggers, exclusions, endorsements, conditions, and claims procedures. For enterprises managing thousands of vendors, that creates a practical problem.

A certificate of insurance confirms that a policy exists. It does not show whether the policy will respond to the vendor failure scenario an enterprise cares about. It does not reveal whether exclusions narrow coverage, whether the enterprise has a direct recovery path, or whether the policy meaningfully addresses ransomware, dependent business interruption, data compromise, regulatory costs, notification expenses, remediation, or third-party claims. Vendors, rightly so, are likely to purchase policies that cover their needs with priority (price is often the primary consideration), and the vendor's insurance broker has a fiduciary duty to pursue cyber coverage consistent with that objective. How, then, does an enterprise client gain real confidence in the adequacy of a vendor's policy without the extraordinary effort of reviewing each vendor policy, in the context of specific loss scenarios.

The reality for most enterprise risk transfer programs is that while the existence of cyber insurance limits are frequently verified, there isn't robust consideration as to whether the coverage is fit for purpose. This puts Risk Managers and CISOs in a difficult position: they are expected to manage vendor exposure, without granular insight into the very risk transference mechanisms on which they depend.

Recovery should not require a legal fight

The third gap is the reliance on litigation, subrogation, indemnity demands, or other adversarial processes after a vendor-caused loss.

When an enterprise suffers loss because of a vendor's breach, control failure, outage, error, or omission, it often must pursue the vendor, the vendor's insurer, or both. That process consumes internal resources, strains vendor relationships, and delays financial recovery. It is slow, expensive, and distracting. It consumes attention when the enterprise should be focused on incident response, remediation, customer communications, regulatory obligations, and business continuity.

Further, if a vendor's policy is exhausted by other claimants, or otherwise doesn't respond, many vendors don't have the balance sheet to cover losses in the absence of viable insurance, again leaving the enterprise client to absorb the loss directly.

Cyber events move quickly. Traditional recovery mechanisms do not. A model that depends on litigation after the damage is done is not a complete solution for operationally critical third-party risk.

What better looks like

Requiring vendors to buy higher limits may help at the margins, but it does not solve the structural problem. Larger shared limits are still... shared. Better cert processes may better illuminate, but don't solve the issue. And the litigation-intensive resolution of vendor-induced loss continues to be inefficient, slow, and problematic, regardless of aggregate limit values or coverage clarity.

The market needs a new approach for third-party cyber risk transfer: dedicated, ring-fenced limits tied to specific enterprise-vendor relationships; standardized policy language for that applies evenly across all vendor-caused cyber events for a given enterprise client; direct recovery for the exposed entity and integration into existing enterprise TPRM workflows.

Further, the market needs to incentivize good cyber hygiene in a more tangible way. Vendors with stronger controls, better incident response practices, clearer contract discipline, and stronger operational resilience should benefit. Risk transfer should reward better security, not simply add another onboarding checkbox.

The next market advantage

Carriers, brokers, reinsurers, enterprises, vendors, and TPRM platforms can keep treating today's model as "good enough," or they can build something better.

Enterprise clients need more than a certificate and the hope that a shared policy structure still has limits available after a widespread loss. They need protection they can confidently evidence, understand, and use when a vendor incident reaches their business.

The next generation of third-party cyber risk transfer will be shaped by those who make vendor risk more transparent, more insurable, and more accountable before the next systemic event forces the issue.

Written By: Josh Ladeau - Trium CEO

Trium Cyber TRAM Services Team: riskservices@triumcyber.com