WHITE PAPER

Third-Party Risk, From a Distance

Assessing any partner without a questionnaire, a deployment, or their participation.

Somewhere in your vendor population, right now, there is a supplier whose network is doing something it shouldn't. You will not learn this from their questionnaire, which they completed carefully and in good faith. You will not learn it from their security rating, which is respectable. In the standard model of third-party risk you will learn it when they disclose, weeks or months from now, and your incident response will begin at the same moment as everyone else's. This paper is about a different model, in which you learn it first, from direct observation, and the vendor conversation begins with evidence instead of a breach notification.

01 The proxy problem

Third-party risk programs run on borrowed evidence, and the two main currencies are both proxies. The questionnaire asks a vendor to describe their controls, and what it measures, with precision, is the vendor's ability to answer questionnaires. The answers are drafted by the people being evaluated, describe intentions as often as implementations, and age from the moment of signature. Nobody involved is lying. The instrument is simply pointed at the wrong thing.

Security ratings improved on this by removing the vendor from the loop, scanning from outside and scoring what is visible: open ports, patch cadence, certificate discipline, leaked credentials. This is real signal about hygiene, and hygiene correlates with outcomes. But a rating describes the quality of the locks. It cannot tell you whether anyone is currently inside the building. An organization can hold an excellent rating while an intruder moves data out of its network, because outside-in scanning does not, and structurally cannot, observe compromise.

The event a third-party risk program exists to catch, a partner in live contact with an adversary, is exactly the event neither instrument sees.

02 The economics of assessment

There is also an arithmetic problem. Meaningful diligence, deep questionnaires, evidence review, on-site assessment, costs enough per vendor that programs ration it. The common shape: the twenty or fifty most critical suppliers get an annual review, a middle tier gets a rating subscription, and the long tail, hundreds or thousands of vendors, gets a contract clause and hope. The rationing is rational. It is also exactly wrong for the threat, because attackers choose targets by access, not by procurement tier. The regional print vendor with a VPN into your environment is a better door than your cloud provider, and no one has assessed the print vendor since onboarding.

Annual cadence compounds the problem. A vendor assessed in March and breached in May carries a clean bill for ten more months. Whatever an annual artifact measures, it is not the vendor's present condition.

03 Observation without participation

Vigilocity's approach removes the constraint that makes rationing necessary: the cost of the subject's participation. Mythic monitors adversary infrastructure at global scale, tracking threat actors as they register domains, stand up command and control, and stage campaigns, then observes which organizations' networks are in live communication with that infrastructure. The collection happens on the adversary's side of the exchange.[1]

Assessing a third party therefore requires nothing from the third party. No agent on their endpoints, no access to their environment, no questionnaire in their inbox, no notification that an assessment is underway. If a supplier's network is talking to infrastructure Mythic tracks, that fact is visible from a distance, on the day it becomes true. If it is not, that is visible too, and it is a statement about observed behavior rather than described controls.

Because per-vendor cost no longer scales with vendor cooperation, the whole population can be watched at once, continuously. The tiering question, which vendors deserve real assessment, dissolves. All of them get it, including the long tail, including the fourth parties you know about, and the assessment never goes stale because it is not an annual artifact. It is a standing observation.

	QUESTIONNAIRE	SECURITY RATING	DIRECT OBSERVATION
What it measures	Described controls	Visible hygiene	Live contact with adversary infrastructure
Vendor participation	Required	None	None
Detects active compromise	No	No	Yes, with evidence
Freshness	Annual, ages immediately	Continuous	Continuous
Population coverage	Top tier only	Broad	Entire population
Output	Attestation	Score	Timestamped evidence

04 What the signal looks like

A finding is not a score movement. It is a record of an observed exchange: this supplier's network, in communication with this adversary infrastructure, beginning at this date and time, over this port and protocol, with this data leaving and this implant responsible. Mythic ranks findings by material impact, what is being taken and from whom, so a program managing a thousand vendors is not triaging a thousand alarms. It is reading a short list of suppliers that matter this week, ordered by consequence, often before the supplier knows anything is wrong.

Evidence changes the vendor conversation more than any contract clause ever has. The traditional escalation, "our rating service flagged you," invites a dispute about methodology. An observed exchange invites incident response. In practice the call is collaborative rather than adversarial: you are handing a partner the first concrete account of their own breach, with enough technical detail, timestamps, addresses, the implant

involved, for their responders to act on immediately. Vendors do not enjoy these calls. They do, reliably, take them seriously.

05 Running a program on observation

Observation does not replace the machinery of a third-party risk program; it feeds that machinery something worth processing. The working pattern among teams using Mythic this way:

01 **Baseline the population.** Put every vendor with network access, data access, or operational dependence under observation, not a tier. The long tail is where the surprises live.

02 **Wire findings to response, not to scoring.** A confirmed exchange is an incident at a partner, and the useful reaction is notification and containment planning, not a risk-register update. Decide in advance who calls the vendor and what evidence they hand over.

03 **Let evidence drive the contract.** Notification obligations, remediation timelines, and audit rights all negotiate differently when findings are observable. Several teams now write continuous monitoring into vendor agreements explicitly.

04 **Keep the questionnaire, demoted.** Attestation still has a place: governance, insurance, regulatory files. It stops being the detection layer, which it never was anyway.

The regulatory current runs the same direction. NIST CSF 2.0 elevated supply chain risk to its own governance category, GV.SC, with an explicit call for monitoring suppliers through their lifecycle rather than at onboarding. NIS2 and DORA in Europe push supervision of critical ICT suppliers toward the continuous and the evidentiary. New York's Part 500 expects third-party policies grounded in the risks vendors actually present.[2] None of these regimes names a technology, but all of them describe capabilities that attestation-based programs strain to demonstrate and observation-based programs simply have.

06 Conclusion

Third-party risk became a paperwork discipline because the alternative did not exist: there was no way to know a partner's actual condition without their cooperation, so programs measured cooperation and called it assurance. That constraint is gone. A partner's live relationship to adversary infrastructure can be observed from a distance, continuously, across the entire population, with evidence attached to every finding. What remains of the old model is worth keeping for what it is, governance. What replaces it at the center of the program is the thing the program was always supposed to produce: knowledge of which suppliers are in trouble, while there is still time to act on it.

07 Sources

1. Vigilocity, "Reverse Attack Surface Analysis: Watching the breach from the attacker's side," 2026. vigilocity.com/resources
2. NIST, "The NIST Cybersecurity Framework (CSF) 2.0," February 2024, GV.SC. · Directive (EU) 2022/2555 (NIS2), Art. 21. · Regulation (EU) 2022/2554 (DORA), Ch. V. · 23 NYCRR Part 500, §500.11.

ABOUT VIGILOCITY

Vigilocity builds technology for identifying and eliminating threat actor infrastructure, drawing on decades of counterintelligence operations inside adversary networks. Its platform, Mythic, confirms material breaches by observing them from the adversary's side.

BRIEFINGS

- **Request a briefing**

vigilocity.com/contact