

Modernizing Diligence Workflows for Scale

AI, Vendor Risk, and National Security: Findings from Quantifind's Stakeholder Advisory Board

Summary

Over the past several years, vendor vetting across the Defense Industrial Base has become a problem of scale, rather than access to pertinent information. Recent incidents, including a foreign-sourced component transmitting data from a [UK naval drone](#), prove that defense supply chain vulnerabilities are not an abstract threat. Simultaneously, acquisition teams face a growing paradox: fewer people, more transactions, and pressure to move faster while working with more vendors in order to leverage as much innovative technology as possible. The story has shifted from having a handful of prime vendors to embracing a chorus of small and emerging businesses to deliver capability faster. AI, particularly when paired with agentic workflows, has emerged as the only viable path to closing the gap between need and capability, but adoption depends on institutions trusting the reasoning behind AI decisioning, clear data provenance, improved communication across government silos, and resolving the tension between moving fast enough to keep pace with adversaries and moving responsibly with powerful new technology.

Key Takeaways

- **Vendor vetting has become a scale problem, not an information problem.** Effective oversight of government vendors must extend to their suppliers and business relationships — and there are more than 200,000 direct suppliers with over 200 million relationships. No manual process, and no incremental increase in headcount, can keep pace with that volume.
- **Converging pressures are straining the system.** A shrinking, overextended workforce is being asked to check more types of risk, move faster, and vet more vendors, while working across siloed offices with different tools, inconsistent risk appetites, and no single owner of FOCI risk.
- **AI is the only viable solution, but capability alone is insufficient.** Purpose-built platforms like Graphyte can screen at scale, investigate deeply, and flag risk continuously, but adoption depends on transparency, consistency, and overcoming cultural resistance.
- **Technology can't fix policy and governance gaps on its own.** Information-sharing restrictions and unresolved questions about where prime contractor liability actually sits will limit impact no matter how good the tooling gets.
- **This is a moving target, and the pace itself is part of the problem.** AI-driven detection tools must evolve fast enough to keep pace with adversaries' tactics, while still working through the broader debate over how quickly AI development itself should move.

More is Different

A weapons program that once meant tracking on the order of 30,000 discrete parts now means tracking roughly 200,000 direct suppliers. Once extended through those suppliers' own vendors, the number of relationships balloons past 200 million. Layered onto a shift from one-time vendor sign-off to continuous, everyday monitoring of ownership changes, sanctions exposure, and financial health, the true scope of what needs to be screened across the Defense Industrial Base is a workload with no historical precedent.

Every policy lever currently under discussion multiplies this workload. In May 2026, the Department of War published a [proposed DFARS rule](#) implementing [Section 847 of the FY2020 NDAA](#), which would extend foreign ownership, control, or influence (FOCI) disclosure and mitigation requirements beyond cleared contractors to unclassified contracts and subcontracts valued above \$5 million. The Department estimates the rule would bring roughly 37,740 contractors and subcontractors into FOCI review, many of them for the first time. Already this year, [Section 851 of the FY2025 NDAA](#) barred the Department of War from contracting with any company whose lobbyists represent a 1260H-listed Chinese military company. Next year, a [companion provision](#) extends 1260H restrictions from direct contracting to indirect business, prohibiting contractors from sourcing components through those companies as well. [Executive Order 14415](#), signed in July 2026, closes a parallel gap: starting January 1, 2027, the Department of War will largely stop issuing routine waivers for critical materials, including rare-earth magnets, tungsten, tantalum, and molybdenum, sourced from China, Russia, Iran, or North Korea, replacing them with a formal mitigation-plan requirement.

Taken together, the expectation going forward is that ownership, and the raw materials and components behind it, must be traceable to origin.

This creates a drastic structural mismatch between how fast the threat surface is growing and how the government and industry have historically staffed and resourced the problem. The dynamic resembles a familiar one in physics: some problems don't just get bigger, they become *qualitatively different* once they cross a certain scale. A phenomenon like superconductivity only appears past a threshold and can't be explained by looking at individual components. A 200-million-relationship problem cannot be solved by hiring a few more contractors to do what worked at 30,000; **it requires a wholly different approach.**

Pressures

Discussion identified nine overarching problems and pressures currently defining the acquisition and vetting landscape:

- 1. **Fewer people.** Hiring has not kept pace, no one expects headcount to grow, and the analysts already on staff are stretched across a growing caseload.
- 2. **More work.** Both Congress and the executive branch are expanding what has to be checked: more suppliers, more tiers, more risk categories.
- 3. **Demand for speed.** There has been sustained pressure to improve efficiency and deploy new capabilities at the speed of relevance, not the speed of bureaucracy.
- 4. **Demand for a broader vendor base.** Agencies are being pushed toward more commercial and startup vendors to accelerate innovation.
- 5. **Fragmented ownership.** No single office owns foreign-influence risk end to end, so findings from one function, such as an intel shop, rarely reach the office positioned to act on them, such as contracting. Agencies and offices also define FOCI differently and set different risk thresholds.
- 6. **Siloed risk tools.** Different offices evaluate the same vendors with different tools and different risk appetites, but those assessments do not travel between offices.
- 7. **The burden on primes is real, but historically implicit.** Large contractors are formally responsible for their own supply chains, though that expectation is rarely enforceable in contract language; the government must perform its own due diligence.
- 8. **Smaller companies want to help.** Early-stage and venture-backed companies want to serve the national security mission, but many cannot see the foreign-influence risk in their own cap tables. When classified or sensitive concerns disqualify them, the government often cannot say why, leaving the company unable to fix the problem and the government without a capability it needed.
- 9. **Cultural resistance.** Existing policy already gives agencies sufficient authority to adopt AI tools for vendor vetting, but the workforce and procurement culture is not yet comfortable delegating first-pass review to AI agents.

Today’s Solutions

Existing AI solutions, like those offered by Quantifind, offer solutions to many of the problems identified in these discussions:

Capability	Value
On-Demand Bulk Screening	Tens of thousands of entities can be screened at once, evaluated against risk factors and data sources relevant to an organization’s mission, and returned within minutes as prioritized risk-tagged findings, enabling small teams to keep pace with growing demand for scale and speed.

Deep Research	Analysts can move directly from screening results into Graphyte Search, an investigative UI, to review annotated source documents and investigate direct and indirect risks , including FOCI, criminal activity, suspicious trade, and international watchlists, quickly uncovering risk hidden behind organizations of any size.
AI Summarization	Large Language Models (LLMs) summarize related findings into concise, auditable insights for human analysts , helping to increase the speed at which cases can be reviewed.
Continuous Monitoring	Automated alerts flag changes in vendor risk profiles , such as new owners, investors, key personnel, or adverse media, replacing one-time checks and random sampling with a simple case management tool.
Agentic Decisioning	Agents automatically resolve false positives with auditable justifications linked to supporting evidence, so analysts spend their time on real risk instead of clearing noise.
ICS 206-01 Compliant Reporting	Export findings into shareable reports that meet intelligence community standards for citing publicly and commercially available information, helping to bridge the communication gap between offices.

Quantifind’s Graphyte suite is purpose-built for this work, not adapted from a general-purpose model. Every finding links back to its original source documentation, and every agentic decision comes with a record of the agent's own reasoning, which is a meaningful differentiator from generic large language models where there is no way to verify why an answer came out the way it did.

Transparency also supports trust building with users and eases some element of cultural friction. An intuitive interface matters for the same reason; a tool the workforce can pick up without months of retraining is itself part of the answer to the adoption hurdle.

Critically, none of these tools are intended to replace an analyst's judgment. Rather, they are designed to remove the rote work, so human attention goes to the work that actually requires it.

Remaining Hurdles & The Road Ahead

No AI tool, however capable, resolves every pressure currently facing this community:

- **Information Sharing Rules:** Better tooling cannot override the rules that govern information sharing. If rules prevent an intelligence office from passing findings to an acquisition office, or the government from passing them to industry, the information cannot be actioned.
- **Reciprocity:** A vendor cleared by one program can still be re-vetted from scratch by the next unless agencies agree to accept each other's findings, which is an administrative choice, not a technical one.
- **Culture:** Even a transparent, well-documented tool still asks an institution to trust it. That trust is built over time, through use. Cultural change is both an individual and institutional force, influenced at least as much by how a tool is deployed and integrated as by the quality of the tool itself.

Solutions to this problem will never be static; the time between an adversary's tactic and an industry's countermeasure, especially visible today in areas like drone warfare, is compressing everywhere, including in how quickly ownership structures, front companies, and financing relationships can be reshuffled to stay ahead of a static rule set. The technology and the data sources behind it must evolve alongside adversaries' tactics to evade detection.

The need for rapid development to keep pace with threats also raises a genuine tension with the broader and entirely legitimate debate over whether frontier AI development itself is moving too quickly for safety to keep up. Both pressures are real, and neither should be waved away to make the other easier.

Conclusion

The scale of this problem is no longer a matter of debate; it's a matter of record. AI-driven tools like Graphyte already close much of the gap between the volume of risk and the capacity to review it, but no platform can substitute the policy and cultural changes still required: shared authority over foreign-influence risk, reciprocity across organizations, and a willingness to let evidence-backed automation inform the analysts who make vetting decisions.. Diligence will keep pace with threats only when technology and governance advance together.