DENEXUS™

JUNE 2026 | NEWSLETTER

OT CYBER RISK INTELLIGENCE NEWSLETTER

• Quantify • Reduce • Transfer

Thought leadership on underwriting-grade OT evidence, AI-driven vulnerability discovery, and the DeRISK platform for quantifying and reducing cyber risk.

denexus.io

CONTENTS

Thought Leadership

Underwriting-grade OT evidence: standardize the minimum

Traceability is the new standard for cyber insurance readiness

Thought Leadership

Underwriting-grade OT evidence: standardize the minimum

Underwriting-grade evidence is not a new concept. What is new is that it can now be defined precisely — because the agentic workflow that produces it has reached the market. DeRISK UWA Agentic, which entered early adopter deployment in May with Chaucer Group as the named reference, moves from submission ingestion to actuarial output in minutes: expected loss, maximum foreseeable loss, full loss exceedance curves, premium indication, a structured insurance program with mandatory binding conditions, and every finding traceable to its source document, pipeline step, and model version.

That traceability is the standard, and it changes what evidence is required to do. Evidence that cannot be traced cannot be audited. Evidence that cannot be audited will not survive claims scrutiny when restoration timelines are contested, will not survive renewal conversations when prior-period assumptions are challenged, and will not earn the comparability that scaled capacity requires.

The discipline this month is about converting that abstraction into a five-artifact minimum that any industrial organization can build, govern, and refresh — without expanding the volume of evidence, and without expanding the maintenance burden beyond what an operating team can realistically carry.

Traceability Pipeline

Evidence → Assumptions → Outputs

Five-Artifact Minimum Evidence Bundle

01	02	03	04	05
Network Boundaries	Access Logs	Restore Test Results	Segmentation Validation	Response Drill Outcomes

What matters this month

Traceability is the standard

Evidence to assumptions to outputs must be versioned, reviewable, and refreshable — not produced once and left to drift. Every assumption that bounds an output should carry a date stamp, a methodology reference, and an owner.

Small and structured beats large and inconsistent.

The minimum evidence bundle is five artifacts: network boundaries current to last verification, access logs covering the past quarter, restore test results with timestamps, segmentation validation through actual traffic testing rather than configuration review alone, and response drill outcomes with documented timing and outcomes.

Controls earn confidence when tested — not attested.

A proof bundle built on outdated test results is not a proof bundle. It is a liability when loss occurs and evidence is scrutinized: the gap between what was documented and what was tested is exactly where claims disputes concentrate.

What you can do now

1. Audit your current evidence bundle against the five-artifact standard: network boundaries, access logs, restore test results, segmentation validation, response drill outcomes. Identify which are missing and which have drifted past their useful date. Drift is governance failure — name it and address it.

2. Assign ownership and set a quarterly refresh cadence per facility. Evidence without a named owner will drift. Evidence without a refresh cadence is a point-in-time document, not a governance artifact. The owner is accountable for currency; the cadence is what makes the role visible.

3. For markets: present evidence that is testable, not just stated. The underwriting workflow now exists to process it. Traceability from evidence to assumption to output is the quality signal that changes pricing conversations — and the absence of that signal is what conservative market terms reflect.

Next month

A practical operating model: Quantify, Reduce, Transfer.

Industry News

02

The final mile: AI can find OT vulnerabilities faster than we can safely fix them

Anthropic expanding Project Glasswing to Critical Infrastructure, alongside OpenAI Trusted Access for Cyber Programme.

Industry News

Anthropic expanding Project Glasswing to Critical Infrastructure, alongside OpenAI Trusted Access for Cyber Programme.

Anthropic has expanded Project Glasswing to roughly 150 organizations across more than 15 countries — power, water, healthcare, communications, and hardware — and OpenAI is moving the same direction with GPT-5.5-Cyber through its Trusted Access for Cyber programme. Frontier AI has changed the economics of vulnerability discovery. For industrial operators, that is both good news and a warning.

150+
Organizations
Participating across 15+ countries in Anthropic Project Glasswing

10,000+
High/Critical Flaws
Identified by Glasswing partners using Claude Mythos Preview

Discovery is no longer the bottleneck. Glasswing partners have already identified more than 10,000 high- and critical-severity flaws using Claude Mythos Preview. But OT vulnerability management was never really about finding problems. It is a waterfall:

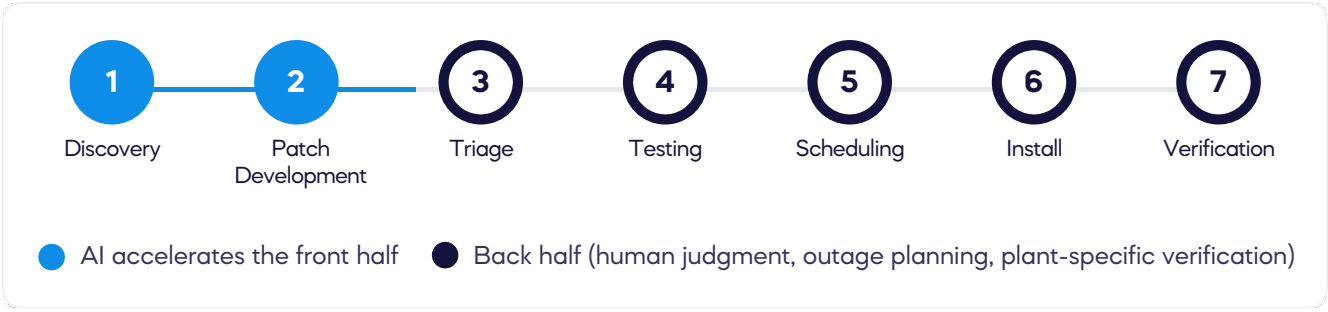

AI accelerates the front half. It cannot replace the engineering judgement, outage planning, vendor accountability, and plant-specific verification in the back half. That is the final mile.

More discovery, same capacity: AI does not create more outage windows, more control engineers, or more vendor support. In real OT telemetry, patches more than 2,000 days old are still outstanding. Pointing faster discovery at an already-saturated system simply grows the backlog — unless prioritization changes.

Read The full article from our Sr Director of OT Cybersecurity – Donovan Tindill

Anthropic expanding Project Glasswing to Critical Infrastructure, alongside OpenAI Trusted Access for Cyber Programme.

The Final Mile Problem: AI Can Find ICS/OT Vulnerabilities Faster Than We Can Safely Fix Them

→ Read More

Triage cannot stop at CVSS. A CVSS 9.8 on an isolated asset with strong compensating controls may matter less than a 7.5 on an exposed remote-access path that reaches multiple plants. The right question is not "how bad is this CVE?" It is "how much financial loss does it contribute, and how much does remediation remove?"

Sourced Platform Simulation Analytics

Directly from the DeNexus DeRISK QVM platform metrics: our continuous simulation engine surfaces the **1–2% of CVEs that drive roughly 90% of real risk** across deployed environments, enabling efficient allocation of mitigation capacity.

The window is short. Anthropic and OpenAI operate under controlled access and safety guardrails. Less-constrained actors — hostile states, criminal ecosystems, loosely governed models — will not. If Western defenders have a temporary head start, the time to turn discovery into risk reduction is now.

How DeNexus Helps

DeRISK QVM ranks vulnerabilities by expected loss reduction, not by severity score — surfacing the 1–2% of CVEs that drive roughly 90% of real risk. It integrates natively with Forescout, Nozomi Networks, Claroty, Tenable, Dragos, Fortinet, and Palo Alto Networks across more than 300 deployments.

1-2%	300+
CVEs Driving 90% Risk	Deployments

Integration Capabilities

DeRISK QVM integrates natively with leading OT security platforms to provide comprehensive vulnerability management and risk quantification.

- Forescout integration
- Nozomi Networks
- Claroty integration
- Tenable integration
- Dragos
- Fortinet
- Palo Alto Networks

DeNexus In June

Where to find us this month

Speaking engagements, webinars, and industry events

Join us at **three key speaking engagements** across virtual and in-person stages. Our team will be hosting two webinars; One on OT cybersecurity & risk quantification, and the other on the future of cyber insurance.

Where to find us this month

Speaking engagements

Fortinet OT Summit
Donovan Tindill • Snr Director of OT Cybersecurity
June8-10 Virtual
→ Learn More

SANS ICS Security Summit
Donovan Tindill • Snr Director of OT Cybersecurity
June 9-11 Orlando, FL
→ Learn More

ISA – OT Cybersecurity
Jose Seara – Founder & CEO
June 16 Prague
→ Learn More

Upcoming live webinars

Two live sessions, each 45 minutes with a live demo, Q&A, and a recorded replay. Presented by Neil Arklie, Donovan Tindill, and Kevin Hamman.

"Industrial OT Cyber Underwriting: From Submission to Binding Decision"

Most industrial OT cyber submissions arrive with IT-derived inputs that miss the exposures that matter. This session walks the full underwriting workflow — submission through to binding decision — and shows where DeRISK UWA Agentic augments the desk with industrial-grade analysis. For cyber underwriters, brokers, and the reinsurance market.

July 2, 2026 60 minutes Neil Arklie, Donovan Tindill, Kevin Hamman

→ Register Now

"Quantified OT Cyber Risk: From Exposure to Reduction"

Quantification is only useful if it changes a decision. This session walks the full path: from industrial OT exposure analysis to financial loss curves with DeRISK CRQ, and from loss curves to prioritised vulnerability remediation with DeRISK QVM — demonstrated on real industrial scope. For CISOs, CFOs, risk officers, and the brokers supporting them.

July 7, 2026 60 minutes Neil Arklie, Donovan Tindill, Kevin Hamman

→ Register Now

The DeRISK Platform — Quantify and Reduce

DeRISK Platform Overview

One simulation core for two critical functions: Quantify and Reduce

The DeRISK Platform covers two of the three moves in the model — **Quantify** and **Reduce** — through CRQ and QVM, built on one simulation core. One quantifies risk in dollars; together they tell you which work removes the most of it.

DeRISK CRQ

Cyber Risk Quantification

DeRISK QVM

Quantified Vulnerability Management

DeRISK UWA Agentic

Agentic Underwriting

DeRISK CRQ — Quantify

Cyber Risk Quantification translates OT exposure into Expected Annual Loss and Value at Risk, then simulates which controls and projects reduce that loss the most per dollar spent — at the facility and across the portfolio. It gives CISOs, CFOs, and boards a defensible financial view of risk instead of a red-amber-green dashboard, and reframes the cyber conversation as a capital allocation conversation.

- Translates OT exposure into financial loss
- Simulates control effectiveness
- Portfolio-level analysis

DeRISK QVM — Reduce

Quantified Vulnerability Management ranks every CVE by expected loss reduction, not by CVSS score. CVEs are matched to your asset inventory and network context, cross-referenced against active ICS-CERT and CISA advisories, mapped to MITRE ATT&CK for Enterprise and ICS through a patent-pending AI pipeline, and run through a simulation-twin of your network that computes the financial loss each vulnerability contributes. The result surfaces the 1–2% of vulnerabilities that drive roughly 90% of real risk — and lets you simulate remediation before committing an outage window.

- Ranks by expected loss reduction
- Cross-references ICS-CERT advisories
- Simulates remediation before outage

Ready to see the DeRISK Platform in action?

Schedule a demo with our team to see how DeRISK can transform your OT cybersecurity strategy.

Book a Demo →

DeRISK UWA Agentic — Transfer

AI-native underwriting platform for the cyber insurance market

DeRISK UWA Agentic completes the Quantify – Reduce – Transfer model. Our AI-native underwriting platform for the cyber insurance market — underwriters, MGAs, and Lloyd's syndicates writing industrial and OT risk — launched in May, entering early-adopter deployment with Chaucer Group as the named reference.

Underwriting OT exposure has always been hard — limited evidence, inconsistent submissions, and not enough specialist capacity to assess every risk in depth. DeRISK UWA Agentic closes that gap. Its agentic workflow moves from submission ingestion to actuarial output in minutes — expected loss, maximum foreseeable loss, full loss exceedance curves, premium indication, and a structured insurance program — with every finding traceable to its source document, pipeline step, and model version.

Agentic Workflow

Moves from submission ingestion to actuarial output in minutes: expected loss, maximum foreseeable loss, full loss exceedance curves, premium indication, and a structured insurance program.

Traceability

Every finding traceable to its source document, pipeline step, and model version. Evidence that cannot be traced cannot be audited.

Augment your underwriting team — without hiring.

→ Learn More

→ Book Demo