INFO~TECH
RESEARCH GROUP

Qualys Enterprise TruRisk Platform

Exposure Management for the Mythos Era



Jon Nelson
Principal Advisory Director,
Info-Tech Research Group

The limited release of Claude Mythos Preview marked a structural shift in the threat environment. Frontier AI can now discover and exploit software vulnerabilities at machine speed, and the volume of disclosed vulnerabilities is set to rise sharply as that capability spreads. Verizon's 2026 DBIR already ranks vulnerability exploitation as the number one initial access vector. The defensive challenge is no longer finding vulnerabilities; it is verifying, prioritizing, and remediating them faster than attackers can act.

The Qualys Enterprise TruRisk Platform is a cloud-native exposure management suite built around that challenge. It consolidates detection across endpoints, cloud, identity, and applications, then applies threat intelligence and business context to narrow millions of findings to the few that represent real risk. Its differentiation in the Mythos era is the remediation engine: patch orchestration, AI-powered patch reliability scoring, and patchless mitigation.

Vulnerability detection is becoming a commodity. The actual value sits in prioritization and machine-speed remediation, and Qualys is one of a small number of vendors building for that future. The platform is a strong fit for organizations consolidating a fragmented stack, running heterogeneous multicloud environments, and ready to give security a path to drive remediation. It is a weaker fit for committed best-of-breed buyers and for those who want detection alone, where its remediation strengths go unused.

Introduction

Two friends are hiking in Africa when one of them notices a lion following. The first friend stops and puts running shoes on. The second says, "You're crazy, you can't outrun a lion!" The friend replies, "I don't have to outrun the lion, I just have to outrun you."

For 30 years, that proverb has been the unspoken reality of enterprise security. The premise was that a determined attacker with sufficient time and resources could breach any system. We didn't have to be perfect; we just had to make it expensive and annoying enough that a rational attacker redirects to a softer target. Network segmentation works because lateral movement gets tedious. Patch prioritization works because an opportunistic attacker who finds you current on the top exploited CVEs moves on to someone who isn't. Even the language of security maturity, the

idea of being “ahead of your peers” or “a hard target,” is the language of relative defense. The whole paradigm rests on the assumption that the attacker is rational, resource-constrained, and willing to give up.

That assumption is now obsolete.

In April 2026, Anthropic unveiled Claude Mythos Preview, a frontier model capable of autonomously discovering and exploiting vulnerabilities in production software at the depth of an experienced human researcher but at machine speed. Anthropic released Mythos to roughly 50 major software vendors under Project Glasswing. Eventually, everyone will have it. In Glasswing’s first month, partners collectively found more than 10,000 high- or critical-severity vulnerabilities, with several reporting bug-finding rates more than ten times higher than before. Anthropic’s May update is clear: The bottleneck in software security has shifted from finding vulnerabilities to verifying, fixing, disclosing, and patching them. The defensive coalition Anthropic assembled buys the industry a head start, not a permanent advantage. The shift is already visible in breach data. Verizon’s 2026 DBIR reports that exploitation of vulnerabilities has become the number one initial access vector, surpassing both credential abuse and phishing for the first time.

The implications are both tactical and strategic. In the short term, attackers holding undisclosed exploits face a value problem. A zero-day is like fruit on a shelf: It has real value, but that value decays as the fruit ripens and eventually spoils. Mythos has just moved the expiration date forward dramatically. Rational holders, whether state actors or criminal sellers, will not wait. They will mark down, sell off, or use what they have before the inventory rots. The next six months will see accelerated use and accelerated liquidation of inventory that would otherwise have been held in reserve.

The longer-term picture is more consequential. What comes after Mythos is not a faster human attacker. It is a tireless one. An agent does not get bored. It does not weigh opportunity cost. It does not pick the easier target because there is no “easier” when parallelism is free. It works every door, every window on every street simultaneously, indefinitely, and the long tail of vulnerabilities we used to dismiss as “unlikely to be exploited in practice” becomes a queue the agent will eventually compromise.

This breaks more than vulnerability management. It breaks the economic logic underneath defense in depth. Deterrence assumes a deterrable adversary. Delay assumes the attacker will eventually leave or that the target will detect them before damage is done. Prioritization, as most organizations practice it, assumes that being slightly less attractive than the next target is a survival strategy. None of those assumptions hold against a Mythos-class attacker, and pretending otherwise is the most dangerous unforced error security leadership can make in the next 12 months.

The era of relative defense is over. What replaces it, and which vendors are building for it, is the subject of this tech note.

Product Overview

Qualys offers a cloud-native security suite organized around the Qualys Enterprise TruRisk Platform, the analytical and AI layer that consolidates exposure data across the environment and produces prioritized, validated risk findings. Collection spans agents on endpoints and servers, scanners for network and web applications, and API-based connections to cloud workloads, identity systems, and AI infrastructure. Qualys also ingests data from third-party sources including SAST, CSPM, EDR, and CMDB tools, allowing buyers to retain existing investments while consolidating analysis in one place.

The product is built around the risk operations center (ROC) concept, which Qualys positions as the proactive counterpart to the SOC. The SOC detects and responds to active threats. The ROC discovers and reduces exposure before threats materialize. The agentic AI layer, including Agents Sid and Nova for continuous discovery and prioritization, and Agent Val for exploit validation, runs across the full collection and analysis stack.

Features and Capabilities

Core Features

The Qualys Enterprise TruRisk Platform consolidates exposure detection and risk analysis across the asset surface. Coverage spans eight capability areas. Exploit

validation and remediation are the key capabilities: They are where vulnerabilities convert into reduced risk and where Qualys has invested most heavily for the Mythos era.

Capability	Coverage
Asset and attack surface discovery	CAASM (cyber asset attack surface management) and EASM (external attack surface management) provide internal and internet-facing asset inventory, including unmanaged assets, shadow IT, and exposed services. Qualys reports that 45% of CAASM customers adopted the product specifically to address inaccurate or incomplete CMDB data.
Vulnerability detection	VMDR (vulnerability management, detection, and response) provides vulnerability scanning across endpoints, servers, network devices, and cloud workloads, augmented by configuration assessment and policy compliance scanning. Qualys is moving detection toward near-real-time through three mechanisms: a version-based fast path that matches agent-reported package inventory against a live vulnerability database without waiting for a traditional signature, continuous per-detection delivery rather than daily batches, and event-driven evaluation triggered by package changes rather than scheduled scans. Qualys states a target of reducing time from CVE publication to customer visibility from the current 12-to-16-hour window toward under 30 minutes. Frontier AI models from Anthropic and OpenAI accelerate the underlying vulnerability research, with findings validated by human experts.
Cloud security	CNAPP capabilities span CSPM (cloud posture), KSPM (Kubernetes posture), CWPP (workload protection), DSPM (data security) and CIEM (cloud identity entitlement management) across AWS, Azure, GCP, and Oracle Cloud.

Capability	Coverage
Web application & API security	TotalAppSec (web application & API scanning) provides DAST (dynamic application security testing) for web applications and APIs, including authenticated scanning and API discovery.
Identity exposure	ITDR (identity threat detection and response) covers identity provider integration (Entra, Okta), permission analysis, and toxic permission combination detection. Coverage of identity-driven attack paths is a relatively recent addition to the suite.
AI security	AI-SPM (AI security posture management), shipping for approximately 15 months, provides discovery and posture assessment for AI models, AI infrastructure, AI software, AI agents, and MCPs, covering training data exposure and application configurations.
Exploit validation	TruConfirm powered by Agent Val autonomously validates whether a discovered exposure is actually exploitable in the customer's specific environment, testing against the real controls and configurations in place rather than assuming a vulnerability is exploitable because it exists. Each validation produces one of three outcomes: exploitable, blocked by a compensating control, or unreachable. EternalBlue is only flagged exploitable if SMBv1 is actually enabled; Log4Shell only if the JNDI path is live. Agent Val drives this continuously, selecting targets, running validation safely in production, and revalidating after remediation. Qualys reports more than 8 million validations in the last 12 months. The architecture is designed for production safety: no data exfiltration, no writes to disk, no persistence, and non-blocking execution.

Capability	Coverage
Remediation	TruRisk Eliminate provides patch and patchless remediation across Windows, macOS, Linux, and third-party applications. Beyond patching, it offers four additional remediation paths: uninstalling bloatware and end-of-life software, applying mitigations for vulnerabilities that cannot be patched (registry changes, service stops, port closures, host-level firewall rules), running custom scripts from a prebuilt library, and isolating assets as a last resort. An AI-powered patch reliability score, built from internet-wide signals and rollback data across Qualys' agent base, predicts deployment risk before patches go out. Wave-based deployment rolls patches out in progressive rings with success thresholds between stages. Qualys reports 150 million patch deployments in the last 12 months, of which 40 million were autonomous (no human intervention), with a rollback rate under 0.1%.

The agentic AI layer cuts across these capabilities. Agent Nova performs continuous discovery and correlation. Agent Val runs validation against discovered exposures. Agent Rocky is a conversational interface that allows security teams to query the platform in natural language and receive prioritized findings.

Differentiating Features

Security-owned remediation as an architectural commitment. Most exposure management vendors stop at prioritization and hand the remediation problem back to IT through a ticket queue. Qualys is one of the few making the explicit argument that patching belongs with security, and the company is investing in the operational infrastructure required to make that argument credible: AI-based patch reliability scoring, wave deployment with progressive rings, and auto-rollback when deployment signals degrade. Whether this works in practice is a question of organizational readiness, not product capability.

A unified collection and analysis surface across the full asset spectrum. Qualys covers endpoints, servers, network devices, cloud workloads, identity systems, web applications, APIs, and AI infrastructure within a single suite, and it ingests exposure data from third-party tools (CSPM, EDR, CMDB, and others) through a large connector library. Several competitors are stronger in specific slices of this surface, particularly cloud and endpoint. What Qualys offers that the point solutions cannot is unified prioritization across the full surface, drawing on both its own sensors and the tools a buyer already runs. For organizations with a heterogeneous attack surface, that unified view has real analytical value, and the platform-as-aggregator model means consolidation does not require ripping out existing investments.

Patchless remediation as a distinct capability for the gap between disclosure and patch. The patch tsunami creates a specific operational problem. More vulnerabilities will be disclosed than can be patched before compromise, and many will have no vendor patch when exploitation begins. Qualys addresses this with a remediation model that does not depend on a patch being available. Patchless mitigations make a vulnerability unexploitable with techniques such as registry changes, service stops, port closures, and host-level firewall rules, deployed through the same mechanism as patches. For vulnerabilities with no patch, end-of-life software, or patches too risky to deploy immediately, this reduces risk until a permanent fix is available. Most of the category leaves this work to customer-written scripts.

Analyst Perspective

Exposure management is being marketed as an evolution of vulnerability management. That framing is incomplete. What the category is asking buyers to undertake is not a tool migration but three simultaneous organizational shifts. Each is necessary; none is sufficient on its own; and the vendor pitch tends to assume all three are easier than they actually are.

The first shift is analytical. Vulnerability management asked what flaws exist on this asset. Exposure management asks what risk this exposure creates in this environment, given what the asset is, what sensitive data is on it, who has access to

it, and what protects it. That shift is necessary, but it depends on a foundation most organizations do not have. Risk-based scoring requires asset metadata, business criticality, data sensitivity, identity context, and ownership records that are accurate and continuously maintained. Without that input layer, "contextual prioritization" devolves to slightly-fancier CVSS sorting. The vendors selling you exposure management in this category assume that foundation exists. The reality is that establishing it is a significant cross-functional program of work, and the platform's value is severely capped until this rework is complete. Qualys eases this shift with AI-assisted metadata population.

The second shift is operational. Traditional vulnerability management operates on scan cycles measured in weeks or months. Exposure management must operate in real time, because the environment changes continuously and, in the Mythos era, so does the threat landscape. A weekly scan is a snapshot of an environment that has already moved on by the time the report renders. Most vendors in the category now claim real-time capability, but the claim covers a spectrum. Continuous discovery of asset and configuration changes is becoming common. Continuous validation of exploitability against active threats is much rarer. Buyers should distinguish between the two and ask vendors to be specific about what is genuinely real time and what isn't.

The third shift is organizational. This is the one no one wants to talk about. The traditional model of security finding vulnerabilities and IT patching them is too slow for a threat environment where exploitation windows are measured in hours or minutes. The category's response has been to argue that patching authority should shift to security, and there is merit in that. But security teams are not built to own production change management, and IT teams are not built to operate at Mythos' tempo. Whichever path an organization chooses, security ownership or tighter coordination, it requires capabilities that neither team currently has. The coordination layer between them, the connective tissue that would let them work together at machine speed, is unbuilt in most enterprises.

Mythos' capability is not just speed. It is the ability to reason about exploit paths, chaining individual weaknesses into actual routes through software code. Anthropic has already demonstrated this publicly, including a forged-certificate exploit chain in wolfSSL. That reasoning capability, applied to source code, produces a stream of

newly discovered vulnerabilities. Applied to a production environment, it produces something exposure management has always reached for but never quite delivered: a risk score that reflects actual exploitability rather than theoretical severity.

Exposure management as a product category is still in its infancy, and frontier AI will disrupt it before it has fully formed. The strategic question for the incumbents is not whether they can build a Mythos-level attack chain engine. The investment levels are outside what any security vendor can absorb. The harder question is whether they can incorporate that kind of reasoning into their risk scoring engines by partnering with a frontier model provider.

The same Mythos dynamic that produces the patch tsunami also produces the demand for risk scoring that actually means something. A list of 10,000 CVEs ordered by CVSS score is not useful. A list of ten exploit paths that actually reach crown jewel assets is. Translating the former into the latter requires path reasoning powered by environment context: data sensitivity, identity, configuration state, network topology, and RTOs. Mature vendors in this category already collect that context. They also already do CVE detection against customer environments. What they have not yet done, and what defines the next era of the category, is to wire frontier-class reasoning into the risk scoring layer so that the output is a prioritized list of actual attack paths, and then orchestrate remediation against those paths at machine speed.

Buyers evaluating this category in 2026 should be asking the vendor directly how they are responding to that shift, and whether the answer is buried in marketing language about agentic AI or grounded in a clear strategic commitment to owning the patch orchestration layer rather than the analytical layer. The vendors who get this right will define the category. The ones who do not will not survive.

Conclusions and Fit Guidance

Strong Fit

Organizations consolidating from a fragmented tool stack. Buyers running separate point tools for vulnerability scanning, cloud posture, identity exposure, and asset

inventory will find significant value in Qualys. Collapsing those functions into one platform with risk-based prioritization is the core of what the platform does well.

Heterogeneous multicloud or hybrid environments. An attack surface spread across multiple clouds, on-premises infrastructure, and containers is where Qualys' breadth really shines. The wider and more varied the environment, the more the unified collection and prioritization layer is worth.

Organizations ready to give security a path to drive remediation at machine speed. Qualys' real differentiation in the Mythos era is the remediation engine: patch orchestration, reliability scoring, and patchless mitigation. Realizing that requires security to have a path to drive or coordinate patching. Buyers with the organizational alignment to act on findings at machine speed get the full benefit. Buyers without it get a very good scanner.

No Fit

Committed best-of-breed buyers. Organizations that have deliberately invested in a category-leading cloud security platform and want depth in that domain over breadth across all of them will find Qualys' consolidation pitch in tension with their strategy. Qualys can ingest from those tools, but the buyer is not looking to consolidate, so the central value proposition does not land.

Buyers who only want detection and prioritization. Qualys' detection and prioritization are strong, and a buyer could adopt the platform for those alone. But an organization that will not operationalize the remediation engine, that wants a scanner and will handle remediation entirely elsewhere, is paying for capabilities it will not use. A narrower, cheaper tool serves that buyer better.

The Key Question

The right question is not whether Qualys does exposure management well. It does. The right question is whether the organization is ready to act on what the platform finds. For buyers who can answer yes, Qualys is one of a small number of vendors building credibly for the Mythos-era threat environment, and it's worth shortlisting on that basis alone.