

Head

THE WEEK'S TOP PICKS

PUBLISHED OCTOBER 23, 2026 · FRIDAYS

// SHARE THIS WEEK

▽ Top 5

▽ Top Pick

EDITIONS tail | grep | head | ls | diff | uniq

The best tooling updates that shipped this week.

▸ // HOW THESE PICKS ARE MADE

SHOW

▼ // THE BRIEF HIDE

What stands out in this week's picks, grouped by what it lets you do. Every tool named links to its pick below.

Agents got governable today. Zenity now sits in front of every Claude Enterprise tool call and denies it before execution, no endpoint agent. Entra issues identities for unmanaged agents and tightens access as risk rises. Separately, GitHub's npm package firewall blocks malicious dependencies at install time, and ArmorCode and Qualys both ship remediation that drafts and verifies the fix.

GOVERN**Put a policy checkpoint in front of what an AI agent actually does**

Zenity's runtime layer renders an allow/deny verdict on each Claude Enterprise tool call via Inference Hooks, with read/write Boundaries classification and prompt-injection taint tracking, so a poisoned context cannot reach a write action; Entra adds Agent ID types for unmanaged agents and graduated risk-driven enforcement. GitLab's beta AI Governance Dashboard and **glab govern** CLI give the audit trail to answer what an agent did and under whose authority.

vim ? help

ity · [Microsoft Entra](#) · [GitLab](#)

PROTECT

Stop a malicious npm dependency before it lands on a developer machine

The public-preview package firewall intercepts npm packages rather than flagging them after install, which is the only point that matters for install-script malware. GA license compliance enforcement and GA agentic autofix close the other end: alerts arrive with a proposed patch instead of a ticket.

[GitHub](#)

PROTECT

Hand remediation the whole loop instead of the ticket

ArmorCode Zero groups findings into a single fix, drafts the pull request and rescans to confirm the issue is actually gone, which removes the triage-to-verification handoff that leaves findings open for weeks. Qualys moves patching to continuous ring-based execution rather than change tickets, so deployment waves progress without a human kicking each one.

[ArmorCode](#) · [Qualys Patch Management](#)

VIEW

ISSUE VIEW

scroll



SCROLL

EXPAND IN PLACE

PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED

EVERYTHING

Every change in this week's picks, in full.

all 24 changes

Do you prefer **scroll**?



// THE PICKS

01 Zenity

AI Security Posture & Discovery

02 SentinelOne

Endpoint Security (EDR/XDR)

03 **GitLab**

SAST / SCA / Code Security

04 **Microsoft Entra**

Identity & Access Management

05 **Qualys Patch Management**

Patch Management

01 Zenity

Zenity Launches Universal Runtime Control for Claude Enterprise

AI Security Posture & Discovery

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

// WHY IT MADE THE LIST

DEPTH 5/5

IMPACT 5/5

NOVELTY 4/5

Holds every Claude tool call for an allow-or-deny verdict before Anthropic executes it, via the `tool_call` Inference Hooks event across all Claude Enterprise surfaces with no endpoint agent — and taints sessions where prompt injection is detected so any follow-on tool call acting on injected instructions is blocked. Tool-provenance controls let policy admit only Anthropic-verified MCP servers.

Zenity shipped a new universal runtime control layer for Claude Enterprise, enforcing pre-execution allow/deny verdicts on every tool call via Anthropic's Inference Hooks, backed by read/write Boundaries classification, tool-provenance restrictions, and prompt-injection taint tracking with full audit logging — all with no endpoint agent required.

↳ WHAT SHIPPED • 4 FEATURES most completely described first ? what's the number?

01 Pre-execution tool-call enforcement for Claude Enterprise

NEW

80

Uses Anthropic's `tool_call` Inference Hooks event to hold every Claude tool call for an allow-or-deny verdict before it executes. Enabled via the `Validate tool calls` configuration, which routes supported tool calls to the organization's Zenity security server for policy evaluation. A single configuration covers claude.ai (web, desktop, mobile), Cowork, Claude Code, Claude for Microsoft 365, Claude in Chrome, Claude's built-in browser, and Claude Tag in Slack — no endpoint agent or browser extension required.

Enable pre-execution tool-call enforcement in Zenity so that every Claude tool call is held for a policy verdict before it runs — no endpoint agent needed.

📌 In the Zenity console, go to Claude Enterprise integration settings and enable `Validate tool calls`. Contact your Zenity account team to activate tool-call enforcement on your tenant.

— Names the hook event, config toggle and exact console path.

Zenity Launches Universal Runtime Contr...

launch-20261019-f198282d

02 Tool-provenance controls in Zenity Boundaries NEW

65

Adds tool-provenance controls that use tool-call metadata to restrict access by provider type — platform, Claude application, third-party MCP server, or client tools such as `bash` and local MCP servers in Claude Code — including the option to allow only Anthropic-verified servers and block unverified or local origins.

Restrict Claude Code's MCP server access to only Anthropic-verified servers organization-wide, blocking calls from unverified or local origins before execution.

 In the Zenity console, go to Zenity Boundaries > Tool Provenance, set `Third-party MCP servers` to `Verified only`, and set unverified/local origins to Prevent — applies across Claude Code without any endpoint hook.

— Names provider types and client tool examples with console path.

Zenity Launches Universal Runtime Contr...

launch-20261019-f198282d

03 Prompt-injection taint tracking and denial audit NEW

65

Detects prompt-injection instructions in content Claude reads, marks the session as tainted, and blocks any subsequent tool call acting on that instruction at the pre-execution checkpoint, surfacing the denial reason in Claude. Denials appear in Claude's Activity Feed with a reference to the associated finding, and the full event chain is available through Zenity's extended Compliance API integration for audit and investigation.

Review a blocked prompt-injection-influenced tool call and its full event chain after a Claude Tag denial in a shared Slack channel.

 In Zenity, open Activity Feed > filter by `Denied`, select the blocked Claude Tag event, and follow the Compliance API audit link to see the injected instruction, the tainted session marker, and the associated tool call parameters.

— Explains tainting mechanism and audit trail, UI review path given.

Zenity Launches Universal Runtime Contr...

launch-20261019-f198282d

04 Read/write classification in Zenity Boundaries NEW

60

Extends Zenity Boundaries to classify tool calls as reads vs. writes/deletes/executes, allowing policies such as permitting CRM reads while blocking deletions, bulk updates, external shares, shell commands, or form submissions — without relying on connector or tool names that vary across integrations.

Enable `Validate tool calls` with a Boundary that allows reads but blocks write/delete operations on shared drives — stopping deletions before they execute with no agent

deployed to endpoints.

 In the Zenity console, go to Zenity Boundaries > New Boundary, set scope to Claude Enterprise, enable **Validate tool calls**, classify the target connector as read-only, and set the write/delete action to Prevent.

– Mechanism clear; UI path given but no raw policy syntax.

Zenity Launches Universal Runtime Contr...

launch-20261019-f198282d

GOOD PICK?



GOOD REASON?



02 SentinelOne

announced · Oct 21, 2026

Endpoint Security (EDR/XDR)

SentinelOne's Singularity platform protects endpoints, cloud, identity and data with AI-driven detection, Purple AI and an agentic SOC layer.

// WHY IT MADE THE LIST

DEPTH 5/5

IMPACT 4/5

NOVELTY 4/5

Singularity ADR discovers, governs, and stops rogue or malicious AI agents at runtime across endpoints, cloud, and SaaS — inventorying every agent, tool, skill, connector, and MCP server with its permissions, and spanning Agentforce, Copilot Studio, ChatGPT, and Claude beyond MCP-gateway coverage. A Unified Storyline ties the agentic chain into the same execution graph as EDR.

SentinelOne launched Singularity ADR to discover, govern and stop rogue or malicious AI agents at runtime, paired it with Singularity ASOC — a new suite bundling AI SIEM, Agentic Investigation, Hyperautomation and an AI Data Pipelines layer with 50+ third-party connectors — and extended Prompt Security into Lenovo's ThinkShield portfolio.

↳ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 Singularity ADR for agentic AI security

NEW

80

Introduces Singularity ADR capabilities for agentic AI security: discovers every AI agent, tool, skill, connector, and MCP server across endpoints, SaaS, and cloud providers — sanctioned or not — along with the permissions each holds and the employees behind it, with multi-platform coverage spanning Salesforce Agentforce, Microsoft Copilot Studio, OpenAI ChatGPT workspace agents, and Anthropic Claude managed agents, going beyond MCP gateway coverage; streams every agentic event as Agent Traces telemetry to Singularity Data Lake for visualization, retention, search, and use in dashboards or detection rules; unifies ADR and EDR behavior in a single execution graph (Unified Storyline) showing the user request, AI assistant, MCP server, and resulting file or process in one chain; delivers real-time agent protection with in-line prevention that blocks prompt injection, jailbreaks, credential exfiltration, and tool poisoning before the action executes, intent drift detection that flags the gap between what an agent was asked to do and what it actually did, and agent containment that stops a single agent mid-action without halting the rest of the workflow; and adds policy and execution governance for AI agents with allow/block controls over which agents run, who uses them, and what they can reach, including approval workflows for high-risk tools, plus agentic investigation and governed response that automatically completes an investigation and executes a response within customer-defined boundaries.

Review the Unified Storyline for an AI agent-initiated incident to see the full chain from user prompt through MCP server to process execution in one graph.

📌 In the Singularity Platform console, go to Incidents > select an ADR-flagged event > open the Storyline view to trace the user request, AI assistant, MCP server, and resulting file or process in a single execution graph.

Contain a single rogue agent mid-action — for example one exfiltrating credentials — without stopping the broader agentic workflow.

📌 In the Singularity Platform console, go to ADR > Agent Inventory > locate the offending agent > select **Contain Agent** to stop its actions while leaving sibling agents in the workflow running.

– Names platforms and scope but no UI or command path. Names Data Lake destination but no access path given. Concrete UI steps and full execution chain described. Three mechanisms named with concrete threats and a containment path. Describes controls and workflows but no concrete surface.

launch-20261021-7d5f1e18

02 Singularity ASOC unified SOC suite **NEW**

75

Introduces **Singularity ASOC**, a new integrated suite on the **Singularity Platform** combining **AI SIEM** with built-in incident and case management (surfacing threats from ingested security data so analysts start from prioritized findings rather than raw logs), **Agentic Investigation** that investigates alerts end-to-end using native and third-party security data to autonomously gather evidence and work toward a verdict, **Hyperautomation** that closes the loop from verdict to remediation by running routine response actions and workflows at machine speed, **AI Data Pipelines** with more than 50 out-of-the-box connectors spanning cloud, identity, email, network, SaaS, data platforms, and vulnerability management — including **AWS**, **Google Cloud**, **Microsoft**, **Okta**, **Cisco**, **Proofpoint**, **Netskope**, **Snowflake**, **Atlassian**, **Oracle**, **Tenable**, **Microsoft Defender XDR**, **Palo Alto Networks Prisma**, **Wiz**, **Darktrace**, **Sophos Central**, and **Check Point**, supporting ingestion via **syslog** and **Kafka** today with **OpenTelemetry** and generic API collection coming, and able to filter and route data to third-party **SIEMs** and storage destinations in addition to **Singularity Data Lake** — plus threat intelligence and **IOC** management, and case management.

– Umbrella naming bundled products, no mechanism of its own. Names 16+ vendors, protocols and routing destinations in detail. States benefit briefly without mechanism or specifics. Short description with no mechanism or named surface. Explains the flow but no named surface or steps.

launch-20261021-2547d6a6

launch-20261021-a0322cde

THINNER COVERAGE BELOW

03 PartnerOne program expansion for AI security **NEW**

58

Opens a Singularity ADR early access program to certified PartnerOne partners, adds an AI Recon Assessment Service for partners to lead AI security conversations with prospects, and expands PartnerOne with new Implementation Engineer certifications covering SentinelOne Prompt Security, endpoint and identity protection, AI SIEM, AI Data Pipelines, and Hyperautomation.

– Names three distinct partner programs and certification scope. [launch-20261021-a0322cde](#)

04 **Prompt Security on Lenovo ThinkShield** NEW 45

Adds SentinelOne Prompt Security to the Lenovo ThinkShield security portfolio, extending existing EDR coverage into AI governance for ThinkShield endpoints; enables teams to identify sanctioned and shadow AI services in use across the workforce and supports detection and response against prompt injection and AI misuse on Lenovo endpoints.

– Names product and portfolio but no concrete steps. [launch-20261021-4cb06267](#)

05 **Singularity Credits consumption currency** NEW 42

Introduces Singularity Credits, a unified consumption currency for AI-powered actions across the Singularity Platform, with administrator-configurable guardrails on AI investigation and automated response usage that bound the scope of every automated action to customer-defined limits.

– Names the guardrail concept but no concrete limits or UI. [launch-20261021-a0322cde](#)

[launch-20261021-2547d6a6](#)

06 **Wayfinder expert-led MDR services** NEW 30

Extends Singularity ASOC with Wayfinder expert-led MDR, threat hunting, and incident response services for teams that want human analysts operating alongside the automated system.

– Names the service but no mechanism or specifics. [launch-20261021-2547d6a6](#)

GOOD PICK?



GOOD REASON?



03 GitLab

Uncover the business logic vulnerabilities SAST scanners miss (2026-10-20)

SAST / SCA / Code Security

// WHY IT MADE THE LIST

DEPTH 4/5

IMPACT 5/5

NOVELTY 4/5

The new Business Logic Security Analyzer runs whole-repository scans for broken object-level authorization, missing authorization and mass assignment — vulnerability classes pattern-based SAST cannot detect — and reports them as first-class results in the merge-request security widget. GitLab Secrets Manager also reaches GA, scoping build-time secrets per job via OpenBao with full audit trails.

GitLab's biggest shipments this window are agentic: an early-access Business Logic Security Analyzer that scans whole repositories and merge requests for vulnerabilities SAST misses, a beta AI Governance Dashboard with new `glab govern` CLI commands for auditing agent activity, and a cluster of GitLab Duo Agent Platform flows covering code review, CI/CD pipeline fixes, reviewer recommendation, and natural-language chat.

WHAT SHIPPED • 7 FEATURES most completely described first ? what's the number?

01 Code Review Flow for automated merge request reviews

NEW SUPPORTED PLATFORM

NEW

96

Assign `@GitLabDuo` as reviewer or comment `/assign_reviewer @GitLabDuo` on a merge request to get inline diff, description, and related-file review comments; an automatic reviews option runs a first-pass review on every new MR without manual assignment. Enabled via `Settings > GitLab Duo > Change configuration > Flow execution` (`Allow flow execution` , `Allow foundational flows`), and configurable at project, group, or instance level with settings cascading down the hierarchy. Runs as a CI/CD job on GitLab-hosted or self-managed runners, supports custom review instructions (language conventions, security practices, path-specific rules) pointed at a template project so standards travel across a group, supports self-hosted models for air-gapped or regulated environments, and includes follow-up review that rechecks prior GitLab Duo Agent Platform findings and verifies fixes. Available on `GitLab.com` , GitLab Dedicated, and GitLab Self-Managed at \$0.25 per review.

Trigger an automated first-pass code review on an open merge request without leaving the comment box.

```
$ /assign_reviewer @GitLabDuo
```

Activate the Duo Agent Platform flows for your group so Code Review Flow and Fix CI/CD Pipeline Flow are available to all members.

Settings > GitLab Duo > Change configuration > Flow execution: enable **Allow flow execution** and **Allow foundational flows** > select Code Review and Fix CI/CD Pipeline > Save changes

– Exact commands, settings path, pricing and deployment scope given

Unblock agentic software delivery with ...

Three developer tasks worth handing to ...

02 AI Governance Dashboard for agent activity auditing

NEW INTEGRATION

NEW

95

The beta AI Governance Dashboard, at AI > Governance > Dashboard tab in a top-level group, unifies visibility into GitLab Duo Agent Platform agents and externally connected agents (Claude Code, Cursor, Codex). Setup uses the **glab govern setup** subcommand (requires **glab** CLI 1.122.0+) and **glab govern setup --agents codex,cursor** to install hooks and capture sessions every 30 minutes, with **glab govern doctor** to verify capture is working. A new **read_agent_artifacts** custom role ability grants dashboard access without requiring Owner on the top-level group. A **Store AI audit events** toggle (Settings > GitLab Duo > Change configuration > Data and privacy) enables persistent recording of every agent action, which can then stream to a SIEM or be downloaded as an immutable session copy for evidence.

Capture Claude Code, Cursor, and Codex sessions so all external agent activity shows up in the AI Governance Dashboard.

```
$ glab govern setup --agents codex,cursor
```

Confirm that agent-capture hooks are installed and sending sessions correctly before handing the dashboard to an auditor.

```
$ glab govern doctor
```

Enable AI audit event storage so agent actions start being recorded — without this, the dashboard shows no history.

 In your top-level group, go to Settings > GitLab Duo, select **Change configuration**, then enable **Store AI audit events** under Data and privacy. Requires Owner or Security Manager role.

– Exact CLI subcommands, role name, and config path all given

AI Governance: One view of agent activi...

03 Business Logic Security Analyzer for whole-repo and MR scans

NEW

79

GitLab Duo Agent Platform adds a Business Logic Security Analyzer running whole-repository SAST scans for broken object-level authorization (CWE-639), missing or improper authorization, and mass assignment — classes pattern-based scanners can't catch. Scans trigger via GitLab Duo Chat or the API, or run automatically on merge requests via an attached scan profile once pipelines pass; findings (CWE ID, severity such as Critical, and file location) appear in the Vulnerability Report and the MR security widget. It runs on Anthropic's Claude Mythos 5.1 (a first for any GitLab capability) plus Claude Opus 5.5, Sonnet 5.5, Sonnet 5, and Sonnet 4.6 as GitLab-managed models via GitLab's AI gateway, in early access for GitLab Ultimate on **GitLab.com**, consuming GitLab Credits scaled to attack-surface size. A companion Security Review Flow checks merge request changes for the same business logic vulnerabilities before they merge.

– Names models, CWEs, and reporting surfaces but trigger path is generic

Uncover the business logic vulnerabilit...

Get every merge request in front of the...

04 Fix CI/CD Pipeline Flow for failing pipelines

NEW

77

Selecting **Fix pipeline with Duo** from a merge request's Overview or Pipelines tab has an agent read pipeline logs and propose inline code suggestions, open a new merge request with a fix, or post next-step comments; the same action is available via **Build > Pipelines** for pipelines not tied to an MR, producing a new merge request with the fix. The **AI > Sessions** view lets you follow the agent's reasoning in real time while it runs.

Send a failing pipeline to the agent for diagnosis and a proposed fix when the pipeline is not attached to a merge request.

 In the GitLab UI: Build > Pipelines > [select failing pipeline] > Fix pipeline with Duo

– Exact UI path given, but limited mechanism beyond that

Three developer tasks worth handing to ...

Get every merge request in front of the...

05 Agentic Chat in GitLab Duo sidebar and IDEs

NEW

68

Agentic Chat opens via `GitLab Duo sidebar > Add new chat` and lets practitioners query merge requests, issues, and pipeline data in natural language; it's available in the GitLab UI, VS Code, JetBrains IDEs, and Visual Studio with the `Agentic` toggle enabled. An AI Resource Library ships ready-made prompts such as `Summarize Recent Changes to MR/Issue` and `Debug Failing Pipeline`.

– Named UI path and prompt examples but no deeper mechanism

`Three developer tasks worth handing to ...`

06 Recommend Reviewers Flow suggests reviewers automatically

NEW

68

Recommend Reviewers Flow suggests the smallest set of human reviewers who can satisfy code-owner and approval rules, factoring in expertise, workload, availability, and timezone. It triggers automatically when a merge request is created, in draft or ready state, surfacing recommendations in a dedicated panel with one-click assignment and an explanation of the reasoning so authors can approve or override.

– Clear behaviour described but no exact UI path or command

`Unblock agentic software delivery with ...` `Get every merge request in front of the...`

THINNER COVERAGE BELOW

07 Custom Flows to compose agentic workflows

NEW

25

Custom Flows lets teams compose Recommend Reviewers Flow alongside other agentic workflows they already run.

– Single-line mention with no mechanism or how-to

`Get every merge request in front of the...`

UPDATED VULNERABILITY CHECK [Container scanning detection for PHP2Shell \(CVE-2026-109618 and CVE-2026-109619\)](#)

NEW INTEGRATION [External Secrets Operator \(Kubernetes\) integration](#)

NEW INTEGRATION [Terraform and OpenTofu data source integration](#)

NEW INTEGRATION [Vault-compatible API for OpenBao and Vault CLIs](#)

[GitLab Secrets Manager: Govern secrets like your code \(2026-10-15\)](#)

GitLab Secrets Manager (GA, 19.5) scopes build-time secrets per job via the `secrets:` keyword in `.gitlab-ci.yml`, encrypts at rest with `AES-256-GCM` via OpenBao, streams full `CRUD` audit events with pipeline/job IDs, and inherits group-level secrets through existing project membership.

Declares secrets in `.gitlab-ci.yml` using the `secrets:` keyword, injecting values only into jobs that declare them and fall within scope (branch, protected-branch flag, `production/*` environment wildcards).

Streams every secret create, read, update, and delete event — including pipeline and job IDs — to audit event destinations and the GitLab audit trail.

Writes secret values to a temporary file inside the job and exposes only the file path, limiting exposure in subprocesses, crash dumps, and telemetry.

NEW INTEGRATION Encrypts secrets at rest with AES-256-GCM via OpenBao, keeping plaintext out of the GitLab application layer.

Issues short-lived tokens bound to each job's identity so jobs fetch values directly from OpenBao.

+5 more

NEW INTEGRATION [External Secrets Operator \(Kubernetes\) integration](#)

NEW INTEGRATION [Terraform and OpenTofu data source integration](#)

PHP2Shell: GitLab helps find and remediate critical PHP flaws (2026-10-22)

UPDATED VULNERABILITY CHECK [Container scanning detection for PHP2Shell \(CVE-2026-109618 and CVE-2026-109619\)](#)

Connect your agents to your software lifecycle with the GitLab MCP server (2026-10-15)

GitLab MCP server is now GA in 19.5, adding tools for branches, semantic code search, pipeline inspection, artifact access, group browsing, wiki lookup, and full Duo Agent Platform session control — all enforced under existing GitLab permissions.

NEW INTEGRATION **NEW SUPPORTED PLATFORM** Adds `add_branch` tool to the GitLab MCP server to start a feature branch from an agent.

Adds `get_repository_file` tool to read a file at any ref without extra API calls.

Adds `semantic_search` tool to find code by meaning, answering 'where is this implemented?' without knowing the file or symbol name.

Adds `get_pipeline` tool that returns a pipeline's status, jobs, downstream pipelines, and artifacts on request.

Adds `list_pipelines` tool to surface patterns across recent pipeline runs.

+9 more

GOOD PICK?



GOOD REASON?



04 Microsoft Entra

changelog · Oct 19, 2026

Identity & Access Management

Microsoft Entra is a cloud-based identity and access management platform that secures user authentication, manages permissions, and controls access to applications and resources.

// WHY IT MADE THE LIST

DEPTH 4/5

IMPACT 4/5

NOVELTY 5/5

Dynamic Conditional Access for agents brings risk-driven, graduated enforcement—reducing permissions, blocking writes while preserving reads, narrowing scope, or isolating agent runtimes—and new Agent ID identity types bring unmanaged and locally discovered agents under policy even without prior Microsoft Security integration.

Microsoft Entra expanded its agent identity capabilities with new Agent ID types for unmanaged agents and introduced Dynamic Conditional Access, a risk-driven graduated enforcement system for agent access.

←→ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Dynamic Conditional Access for agents

NEW

72

Entra introduces Dynamic Conditional Access for agents, enabling risk-driven, graduated policy responses that can reduce permissions, block writes while preserving reads, narrow data/resource scope, restrict network paths, shorten or revoke time-bound access, require human approval, or isolate/stop agent runtimes — coordinated across multiple enforcement points including signals from Microsoft Sentinel, Entra, Intune, Purview, Global Secure Access, and Defender.

– Describes graduated enforcement mechanisms but no config keys or commands

2026-10-19-a16a71ec

THINNER COVERAGE BELOW

02 Agent ID coverage for unmanaged and local agents

IMPROVED

45

Microsoft Entra Agent ID now supports additional identity types for agent patterns, including local agents and previously unknown agents discovered through security sensors, so unmanaged agents become policy targets even without prior Microsoft Security integration.

– Explains scope expansion but no config or API specifics

2026-10-19-a16a71ec

GOOD PICK?



GOOD REASON?



05 Qualys Patch Management

changelog · Oct 14, 2026

Patch Management

Qualys Patch Management is a cloud-based tool that automates the discovery, prioritization, and deployment of patches across IT assets to reduce vulnerabilities.

// WHY IT MADE THE LIST

DEPTH 5/5

IMPACT 4/5

NOVELTY 4/5

TruRisk Eliminate 2.0 replaces ticket-driven patching with autonomous RemOps: TruConfidence picks the safest action per validated exposure — patch, mitigation, uninstall, or existing control — and Agent Sara executes it in customer-defined waves and progressive rings with automatic rollback on anomaly. Scanner-based remediation now reaches edge and perimeter devices agent-only tools miss.

Qualys launches TruRisk Eliminate 2.0, an Autonomous RemOps platform built around a new TruConfidence AI orchestration engine and the Agent Sara autonomous remediation agent, moving patch management from ticket-driven workflows to continuous, ring-based execution.

L--> WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Agent Sara autonomous remediation agent NEW 56

Agent Sara is Qualys' autonomous AI agent that executes TruConfidence remediation plans in customer-defined waves and progressive rings, with automatic rollback triggered on anomaly detection.

- Describes rollout and rollback mechanism, no UI or config path given

2026-10-14-af38b70c

02 TruConfidence remediation and orchestration intelligence engine NEW 53

TruConfidence is an AI-powered Remediation and Orchestration Intelligence engine that selects the safest remediation action — patch, mitigation, uninstall, or an existing control — for each validated exposure.

- Explains decision logic and the four action types, no runnable step 2026-10-14-af38b70c

03 Zero-touch and ring-based deployment with rollback monitoring NEW 50

Supports zero-touch automation for high-reliability patches on low-impact applications alongside ring-based deployment for higher-risk changes, backed by operational health monitoring and a sub-0.1% rollback rate.

– Gives concrete rollback-rate figure but no setup steps 2026-10-14-af38b70c

04 **TruRisk Eliminate 2.0 Autonomous RemOps platform** NEW

42

TruRisk Eliminate 2.0 replaces ticket-driven patching with a continuous Decide → Execute → Prove cycle that runs across the full attack surface.

.

– Names the three-stage cycle but no concrete usage path 2026-10-14-af38b70c

05 **Scanner-based remediation for edge and perimeter devices**

IMPROVED

35

Remediation is extended to edge and perimeter devices using existing Qualys scanners, closing coverage gaps that agent-only approaches miss.

– States the extension and rationale but no mechanism detail 2026-10-14-af38b70c

GOOD PICK?



GOOD REASON?



// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · October 23, 2026

The full firehose — every feature release, every day — is in [the newsletter](#).

▷ **Subscribe**