

Head

THE WEEK'S TOP PICKS

PUBLISHED OCTOBER 16, 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.

IDENTIFY**See the chain, not the finding**

Token Security's Attack Path Analysis stitches identity and permission relationships across AWS, Azure, Okta, Entra ID and GitHub into visual escalation graphs, so the work of manually correlating a low-severity role in one cloud with a GitHub token in another gets done for you. Censys adds post-quantum cryptography visibility to host and web scan data, which turns the eventual PQC migration from a guessing exercise into an inventory you can actually sort.

Token Security · Censys

PROTECT**Stop handing out standing credentials and open tool access**

JumpCloud PKI runs a cloud-hosted private CA with automated SCEP/MDM delivery, so EAP-TLS Wi-Fi and VPN/mTLS auth can replace shared PSKs and passwords without anyone standing up and babysitting a CA. Aembit's MCP Tool Access Control applies least privilege at the tool

level through its MCP Identity Gateway, so an agent with a legitimate server connection no longer inherits every tool that server exposes.

[JumpCloud](#) · [Aembit](#)

PROTECT

Keep agent-written code and agent-driven payments from going unchecked

Endor Labs adds AI Security Code Review for architectural risk in PRs, MCP trust scoring, and Package Firewall coverage for VS Code extensions — a dependency surface most teams never scanned at all. Operant AI's Semantic Firewall flags prompt injection and task drift in agentic payment flows across x402, MPP and AP2, where the failure mode is an agent that transacts correctly but for the attacker's goal.

[Endor Labs](#) · [Operant AI](#)

VIEW

ISSUE VIEW

scroll

EXPAND IN PLACE

PREVIEW PANE

INFORMATION DENSITY

everything

BRIEFED

EVERYTHING

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

all 15 changes

Do you prefer scroll?  

// THE PICKS

01	Operant AI	AI Runtime Defense & Guardrails
02	Belkasoft	Digital Forensics & Incident Response
03	Token Security	Non-Human & Agent Identity
04	Endor Labs	Software Supply Chain Security
05	Aembit	Non-Human & Agent Identity

01 Operant AI

Stop Dollar Exfiltration: Operant Agent Protector now defends agentic commerce across the

AI Runtime Defense & Guardrails

Operant AI discovers, detects and defends AI applications, agents, MCP servers, APIs and Kubernetes workloads at runtime.

// WHY IT MADE THE LIST

DEPTH 4/5

IMPACT 5/5

NOVELTY 5/5

Agentic Commerce Defenses gate every agent payment across the x402, MPP, and AP2 protocols — blocking out-of-task spend, enforcing approved wallet destinations, and reserving budget atomically so concurrent agents can't double-spend — covering an attack surface no other tool inspected.

Operant AI shipped Agentic Commerce Defenses for agentic payments, covering the x402, MPP, and AP2 protocols, alongside a new Semantic Firewall that flags payment-related prompt injection and task drift.

→ WHAT SHIPPED • 6 FEATURES most completely described first ? what's the number?

01 Purse-string spending limits with atomic reservation NEW

75

Spending limits can be applied by agent, service, MCP server, user, team, network, and asset, with hourly, daily, weekly, or monthly windows; budget is reserved atomically before forwarding so concurrent agents cannot double-spend the same allowance.

– Names concrete time windows and the double-spend mechanism

Stop Dollar Exfiltration: Operant Agent...

02 Agentic Commerce Defenses across x402, MPP, AP2 NEW

65

Operant Agent Protector now evaluates each agent payment attempt — combining task context, the paid tool call, MCP server identity, payment destination, amount, and org policy — before forwarding it, across the x402, MPP, and AP2 payment protocols.

– Names three protocols and evaluation inputs but no concrete config surface

Stop Dollar Exfiltration: Operant Agent...

03 Commercial scope controls for payment destinations NEW

60

Organizations can define approved MCP servers, payment networks, assets, and wallet destinations; payments to unapproved destinations are blocked regardless of credential validity.

– Names the control surfaces but no API or UI path given

Stop Dollar Exfiltration: Operant Agent...

THINNER COVERAGE BELOW

04 **Operant Semantic Firewall for payment anomalies** **NEW**

55

The Operant Semantic Firewall evaluates paid tool activity against the session's original task to flag prompt injection, task drift, or payment loops even when individual transactions are properly authorized and under budget.

– Describes detection scope but no interface to act on it

Stop Dollar Exfiltration: Operant Agent...

05 **Task-aware payment gating** **NEW**

50

Blocks a payment when the paid tool call falls outside the agent's assigned task, even when credentials are valid and the amount is within budget.

– Clear before/after behaviour but no named config or command

Stop Dollar Exfiltration: Operant Agent...

06 **Audit logging for payment decisions without custody** **NEW**

50

Logs payment requests, policy decisions, submissions, and reported outcomes without storing payment proofs or signing transactions, so Operant never holds custody of customer wallet keys.

– States what is logged and a custody guarantee, no retrieval path

Stop Dollar Exfiltration: Operant Agent...

GOOD PICK?



GOOD REASON?



02 Belkasoft

announced · Oct 9, 2026

Digital Forensics & Incident Response

// WHY IT MADE THE LIST

DEPTH 5/5

IMPACT 4/5

NOVELTY 4/5

Belkasoft X now does physical MTK-dump acquisition over BootROM mode for 23 MediaTek chipsets, decrypts FBE-protected userdata, and runs offline brute-force passcode recovery — a complete path to imaging and unlocking encrypted Android devices that investigators previously could not open.

Belkasoft X adds MTK dump acquisition via BootROM mode: physical imaging, FBE userdata decryption, and offline brute-force passcode recovery for MediaTek Android devices.

COVERAGE

NEW supported platform [MediaTek chipset physical acquisition support](#)

[All of Belkasoft's coverage changes →](#)

- › Adds **MTK dump** acquisition method using BootROM (BROM) mode to create a physical image and extract cryptographic material from supported MediaTek-based Android devices.
- › Adds decryption of the **userdata** partition (file-based encryption / FBE) using cryptographic material obtained during acquisition.
- › Adds offline brute-force passcode recovery via the **Brute-force options** dialog: built-in Belkasoft dictionary (4- and 6-digit numeric), sequential numeric brute-force, and a custom **User dictionary** (.txt, one entry per line) supporting alphanumeric passcodes — running at ~60 candidates/second without device lock-screen limits.
- › Supports importing MTKClient-acquired dumps as a **Mobile image** data source after key-format conversion.

GOOD PICK?



GOOD REASON?



03 Token Security

announced · Oct 13, 2026

Non-Human & Agent Identity

Token Security is a blockchain security analysis platform that audits smart contracts and tokens for vulnerabilities, scams, and security risks.

// WHY IT MADE THE LIST

DEPTH 4/5

IMPACT 4/5

NOVELTY 4/5

Token Security's Attack Path Analysis traces multi-hop privilege-escalation chains across AWS, Azure, Okta, Entra ID and GitHub — modeling human users, service principals, AI agents and GitHub apps as nodes with federated SSO and cross-account role assumption as edges, surfacing only the exploitable paths with the findings that make each hop work.

Token Security launched Attack Path Analysis, mapping multi-hop privilege escalation chains across AWS, Azure, Okta, Entra ID, and GitHub into prioritized, visual attack graphs.

←→ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

01 Attack Path Analysis across cloud and identity providers NEW

80

Models IAM roles, service principals, human users, AI agents, and GitHub apps as nodes in a directed graph, with federated SSO and cross-account role assumption as edges, tracing IAM assume-role and federation chains across AWS, Azure, Okta, Entra ID, and GitHub to surface exploitable paths out of billions of theoretical ones. Results appear as a prioritized, filterable table showing source identity, destination resource, severity, hop count, and environments involved; each path gets a severity score based on destination criticality, source identity risk, administrative access involvement, and exploitability from existing findings. An interactive graph view labels identity nodes by type and environment, shows attacker actions like federated SSO and cross-account role chaining on edges, and a sidebar exposes findings — unrotated credentials, excessive permissions, missing MFA — that make each hop exploitable; a stats bar summarizes total paths by severity and the identities most frequently acting as dangerous starting points or critical targets.

Investigate which single identity fix eliminates the most critical breach paths across your AWS and GitHub environments.

📌 In the Token Security platform, navigate to Attack Paths. Filter the prioritized table by severity **Critical**. Review the stats bar to identify which source identity appears most frequently as a dangerous starting point. Select a path to open the graph view and use the sidebar to inspect the findings — such as unrotated

credentials or missing MFA — that make each hop exploitable, then remediate at that node to collapse the maximum number of paths at once.

– Mechanism, scoring factors, and node/edge model are detailed; UI path given, no API/flag.

launch-20261013-8fe27d22

GOOD PICK?



GOOD REASON?



04 Endor Labs

Introducing AI Security Code Review | Blog | Endor Labs (2026-10-12)

Software Supply Chain Security

Endor Labs is a software supply chain security platform that detects vulnerabilities and malicious dependencies in open source code.

// WHY IT MADE THE LIST

DEPTH 5/5

IMPACT 5/5

NOVELTY 4/5

AI Security Code Review deploys multiple agents to catch auth gaps, CSRF breaks, and crypto misuse that rule-based SAST/SCA misses on every pull request — paired this week with a new CI/CD Security line (pipeline discovery, repository posture management) and AI-generated Zero-Day Patches for novel open-source vulns.

Endor Labs shipped a wide set of agentic-era capabilities this window: a Grok Build integration (Agent Kit, Coding Agent Governance, MCP trust scoring), a new AI Security Code Review for architectural PR risk, full C language support in AI SAST, Package Firewall protection extended to VS Code extensions, and founding membership in the Linux Foundation's Project Akrites for zero-day patch coordination.

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

01 AI Security Code Review for pull requests

NEW

93

Adds AI Security Code Review, which deploys multiple AI agents (Developer, Architect, Security) to analyze every pull request for design flaws and architectural security changes missed by rule-based SAST/SCA tools. It detects security-impactful PR categories including API endpoint exposure, access control changes, authentication modifications, CI/CD workflow additions, cryptographic library swaps, database schema changes, dependency privilege escalation, input validation gaps, memory management changes, network security control modifications, payment processing integrations, and PII data handling; assigns each a priority level (Critical, High, Medium, Low) and surfaces findings in a dedicated Security Review section within each project. It integrates with the Endor Labs GitHub App to auto-trigger analysis on new PRs across a GitHub organization (surfacing submitter and reviewer identity) and connects to Endor Labs' policy engine and existing SCM/project management systems to automate ticket creation and routing, e.g. sending PII findings to GRC teams or infrastructure changes to DevOps.

Enable AI Security Code Review for your GitHub organization so every new PR is automatically analyzed for architectural security risks.

📌 In the Endor Labs console, go to your project settings and enable the Endor Labs GitHub App for your GitHub organization. Once enabled, AI Security Code Review will automatically run on all new pull requests and surface findings in the **Security Review** section of each project.

– Full mechanism, named categories, and exact enablement steps given

Introducing AI Security Code Review | B...

02 Endor Labs Agent Kit for Grok Build NEW 90

Installable via the marketplace or by typing `/add-plugin endorlabs` in the agent chat, Agent Kit brings AI SAST, SCA, and remediation agents into the IDE session. It includes specialized AURI remediation agents — an AI SAST Remediation agent and an SCA Remediation agent — that pull findings from scans, propose fixes, and open pull requests with developer approval, and can run outside the IDE in CI or cloud runs through the Grok Build SDK.

Bootstrap Endor Labs security scanning inside an active Grok Build agent session without leaving the IDE.

```
$ /add-plugin endorlabs
```

– Names exact install command and remediation agents

Endor Labs + SpaceXSI: Securing every s...

03 Package Firewall and detection extended to VS Code extensions

NEW 80

Extends Threat Center malicious-package detection to the Visual Studio Marketplace, covering VS Code extensions alongside npm, PyPI, Go, Maven, and NuGet, and extends Package Firewall blocking to VS Code extensions, preventing unsafe versions from appearing in the marketplace and blocking unsafe extensions installed via the command line. Adds three Package Firewall policy controls for VS Code extensions — **Malware** (block versions identified as malicious), **Minimum package age** (hold back new versions until vetted), and **Exceptions** (always-allow list for trusted extensions) — each settable to block or warn, logs every extension install, and deploys via existing MDM with no agent, wrapper, or certificate required on developer devices.

– Names all three policy controls and deployment model, no exact command

Learn how Endor Labs detects malicious ...

THINNER COVERAGE BELOW

04 AI SAST and SCA support for C NEW 58

Adds full C language support to AI SAST with buildless scanning in the IDE and on PRs, paired with C SCA in the same platform so first-party C code and open-source dependencies are covered under a single analysis run. Integrates C scanning into AURI by Endor Labs, providing an independent security harness for agentic SDLC workflows with function-level call paths, reproducible evidence, and policy enforcement across any agent, model, or CI stage.

– Describes scope and mechanism but no config or command shown

Why Endor Labs AI SAST for C finds what...

05 **Coding Agent Governance hooks in Grok Build** NEW 53

Records every MCP tool call, shell command, and file operation and enforces policies deterministically before each action runs, without relying on a model's judgment. Also blocks confirmed malicious packages at install time during agent-driven dependency installation.

– Explains enforcement mechanism but no config surface named

Endor Labs + SpaceXSI: Securing every s...

06 **Zero-Day Patches and Project Akrites membership** NEW 50

Joins the Linux Foundation's Project Akrites as a founding member alongside Anthropic, OpenAI, Amazon, Chainguard, Cisco, Citi, IBM, JPMorgan Chase, and Microsoft, routing coordinated vulnerability disclosure and remediation through a shared channel. Launches Zero-Day Patches, delivering AI-generated, verified fixes for novel open source vulnerabilities in under 24 hours, and applies function-level reachability analysis across direct and transitive dependencies to rank which zero-day findings are actually reachable from a customer's specific application code.

– Names coalition and mechanism but no user-facing action

Open source carries the world. Patching...

07 **Endor Scores for MCP servers and skills** NEW 38

Rates each MCP server and skill for trust, mapped to the OWASP Top 10 for LLM Applications, giving security teams a basis to approve or block agents and skills.

– Describes scoring purpose but no scale, thresholds or UI path

Endor Labs + SpaceXSI: Securing every s...

Introducing CI/CD Security with Endor Labs | Endor Labs (2026-10-16)

Endor Labs expands its SSCS platform with CI/CD security: pipeline discovery, repository posture management, and artifact signing with provenance.

COVERAGE

NEW detection

Repository Security Posture Management (RSPM) out-of-the-box policies

- NEW** integration Pipeline Discovery — CI/CD tool and SCM visibility
- NEW** compliance framework

SLSA, NIST SSDF, SOC2, PCI alignment via RSPM and artifact signing

All of Endor Labs's coverage changes →

- › Adds Pipeline Discovery to map every tool touching code across repositories and CI/CD workflows, surfacing shadow engineering, security coverage gaps, and policy violations.
- › Adds Build Integrity Verification via artifact signing, attaching provenance information (source repository, branch, commit, repository owner) to signed packages and containers.

Container Scanning + SCA = Better Together | Endor Labs (2026-10-16)

Endor Labs adds container image signing at build time, signature verification at deploy time, and unified cross-dependency PR comments for vulnerability findings.

COVERAGE

- NEW** vulnerability check
- Container image OS and application package vulnerability scanning
- NEW** integration
- Runtime CNAPP integration for container tracing and admission control
- All of Endor Labs's coverage changes →

- › Adds container image signing at build time and signature verification at deploy time, with scan results consumable by CNAPP and admission controllers.
- › Includes container vulnerability findings in pull-request comments alongside application dependency findings, so developers see all impacted packages in one place.
- › Enables tracing of a production container finding back to its source repo, release artifact, and build via integration with runtime CNAPP tools.

Your model, our analysis: BYOK for Endor Labs AI SAST | Endor Labs (2026-10-14)

BYOK routes AI SAST and AI Chat model requests through your own LLM provider account with an encrypted key store, provider allowlist, and tenant isolation via the Endor Labs AI Gateway.

COVERAGE

- NEW** supported platform First-class C language support for AI SAST

NEW integration [BYOK provider integrations for AI SAST and AI Chat](#)

[All of Endor Labs's coverage changes →](#)

- › Adds bring your own key (BYOK) support for AI SAST and AI Chat, routing all model requests through your own OpenAI, Azure OpenAI, Gemini, or Vertex AI account instead of Endor Labs' provider accounts.
- › Routes BYOK requests through the Endor Labs AI Gateway with an encrypted key store, provider allowlist, and tenant isolation.
- › Supports passing provider API keys as environment variables during AI SAST scans so keys are never stored with Endor Labs.
- › Honors the region you have configured in Azure OpenAI or Vertex AI, so model requests run in your chosen data-residency region.
- › Adds custom organizational context ingestion from Infrastructure as Code files, repository files such as `AGENTS.md` and `CLAUDE.md`, triage feedback, and uploaded threat models to teach AI SAST how your application is deployed.

Endor Labs and Bright Security Partner to Close the Gap Between Discovery and Remediation

Endor Labs and Bright Security integrate reachability SAST with runtime DAST via SARIF handoff, enabling a CI gate and AI-agent-ready evidence for confirmed-exploitable findings only.

COVERAGE

NEW integration [Endor Labs + Bright Security \(Bright STAR\)](#)

[All of Endor Labs's coverage changes →](#)

- › Adds a SARIF-based integration pipeline where Endor Labs' function-level reachability analysis passes a narrowed candidate list to Bright STAR for runtime exploit confirmation, replacing full-surface DAST sweeps.
- › Introduces Endor SAST + Bright STAR as a combined CI gate, catching reachable and runtime-confirmed exploitable code before it reaches the production branch.
- › Enables runtime-confirmed prioritization: findings Bright Security confirms exploitable at runtime are automatically elevated above reachable-but-unproven and unreachable findings.
- › Delivers machine-readable, verified evidence (call path from Endor Labs plus confirmed exploit path from Bright Security) to AI remediation agents, reducing token overhead for reachability and exploitability re-derivation.

Endor Labs Agent Kit brings AURI remediation agents (AI SAST and SCA) into Grok Build sessions and CI, with per-change approval gates and audit-ready evidence.

COVERAGE

NEW integration [Endor Labs + SpaceXAI Grok Build](#)

NEW detections

[Coding Agent Governance](#) — MCP tool call, shell command, and file operation hooks,
[Malicious package blocking at install time via Coding Agent Governance hooks](#)

[All of Endor Labs's coverage changes →](#)

- › Adds **Endor Labs Agent Kit** for Grok Build, installable via `/add-plugin endorlabs` in the agent chat or from the marketplace, bringing specialized AURI remediation agents directly into the IDE and CI/cloud runs through the Grok Build SDK.
- › Adds **AI SAST Remediation** agent (part of the Endor Labs Agent Kit) that traces data flow from untrusted input to vulnerable calls, proposes fixes and test plans, and can open a pull request upon developer approval.
- › Adds **SCA Remediation** agent (part of the Endor Labs Agent Kit) that queries scan results, checks package usage, applies upgrades or Endor Patches, adds tests, and rescans to confirm no new vulnerabilities were introduced.
- › Adds **Endor Scores** rating system for MCP servers and skills, mapped to the OWASP Top 10 for LLM Applications, giving security teams a trust signal for every agent component in use.

GOOD PICK?



GOOD REASON?



05 Aembit

Introducing MCP Tool Filtering: Fine-Grained, Least-Privilege Authorization for AI Agents

Non-Human & Agent Identity

// WHY IT MADE THE LIST

DEPTH 4/5

IMPACT 4/5

NOVELTY 5/5

MCP Tool Access Control enforces per-tool allow/block policies with glob patterns across both tools/list discovery and tools/call invocation, scoped by blended user, group, and workload identity and failing closed — least-privilege for AI agents sharing an MCP server that teams could not express before.

Aembit introduced MCP Tool Access Control, giving the Aembit MCP Identity Gateway fine-grained, least-privilege authorization over which AI agent tools can be discovered and invoked.

→ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 MCP Tool Access Control in Identity Gateway NEW 94

Adds per-tool allow/block rules inside an access policy, independently controlling `tools/list` discovery visibility and `tools/call` invocation. Rules support exact tool names and glob patterns (e.g., `get_user*`, `admin_*`, `list_*`, `delete_*`) so new tools inherit rules automatically, and can be enforced in `Allow` mode (deny unmatched by default) or `Block` mode (permit unmatched by default), selectable per policy. Rules are scoped by blended identity — user identity, group membership, and workload/agent identity — so different principals sharing an MCP server get different tool sets, and the policy is evaluated before deeper inspection layers like CrowdStrike AIDR. The gateway fails closed by default: a missing or malformed tool-filtering policy returns an empty `tools/list` and blocks `tools/call`. Every filtered `tools/list` response or blocked `tools/call` request is recorded via the `application.mcp.contentSecurity.toolFilter` audit event, including which rule matched and why.

Lock a read-only customer-support agent to only `list_*` tools on a GitHub MCP server, blocking all write and delete tools by glob pattern.

📌 In the Aembit console, open your access policy for the customer-support agent, enable MCP Tool Access Control, set mode to `Allow`, and add the rule: Tool pattern: `list_*` Leave all other tools unmatched — they are denied automatically. Save the policy; no MCP client or server changes are needed.

— Names policy fields, modes, patterns, audit event and exact UI steps.

Introducing MCP Tool Filtering: Fine-Gr...

GOOD PICK?



GOOD REASON?



// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · October 16, 2026

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

▷ **Subscribe**