THE CYBER TOOLCHAIN NO. 010

Head

THE WEEK'S TOP PICKS

PUBLISHED SEPTEMBER 4, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

// HOW THESE PICKS ARE MADE

Every feature release from the 907 tools on our watchlist goes into the daily newsletter. Once a week we read the whole field side by side and choose our top picks, judged on two questions:**how deep and complete is the single best capability in the release**, and**how much it changes what you can actually do**.

A tool is judged on everything it shipped that week, so a project that releases daily gets credit for the sum — and still only takes one slot. We would rather run a short list than a padded one.

// THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 Opus 5

A read across the whole week before you read any of it: what stands out in this week's picks, grouped by what it lets you do. Every tool named links to its pick below.

Three unrelated things stand out today: Endor Labs can stop bad packages before install, Cycode is aiming code scanning at authorization and business-logic flaws, and NodeZero added WebApp and Azure pentest coverage so more exposure can be exercised, not just inventoried.

PROTECT

Block risky dependencies before they ever land on a developer machine

Package Firewall enforces malware, vulnerability, and compliance policy during package install, before the artifact reaches disk. That moves the control point earlier than repo scanning or post-build findings, cutting off typosquats and known-bad dependencies before they become cleanup work.

Endor Labs

PROTECT

Test for flaws that do not fit normal scanner rules

Cycode's agentic scanning targets authorization and logic-level issues that are hard to express as static rules, while Edgescan Atomic builds attack plans that include IDOR, business-logic, and multi-identity authenticated testing. The useful change is less hand-authored test design for classes of bugs that usually require human reasoning to exercise.

Cycode · Edgescan

PROTECT

Exercise more of the attack surface with automated pentest paths

NodeZero added WebApp and Azure Cloud pentest operation types plus an Azure Entra ID AD Password Audit; adscan added cross-forest attack path analysis with automated escalation routes and evidence blocks. Teams can validate reachable identity and cloud paths with less manual chaining, and get more defensible proof of how an exposure can be used.

Horizon3 NodeZero · adscan

RESPOND

Move more forensic review offline without losing AI assistance

Belkasoft X now uses BelkaGPT for offline evidence translation and SQLite database classification, alongside custom SQLite artifact mapping and offline map support. That keeps sensitive case material out of external services while reducing manual review of unfamiliar databases, foreign-language content, and mobile artifacts.

Belkasoft

Does **GPT-5.5** read better? 👍 👎

// THE PICKS

01 Cycode SAST / SCA / Code Security

02	**Horizon3 NodeZero**	Exploitation & C2
03	**Endor Labs**	SAST / SCA / Code Security
04	**adscan**	Exploitation & C2
05	**Binary Ninja**	Reverse Engineering

01 Cycode

2 RELEASES · 2026-09-01

SAST / SCA / Code Security

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Agentic Code Scanning adds an agent layer that plans from a threat model and disproves findings before surfacing them, catching logic-level flaws like broken access control that rule engines structurally cannot express — it flagged real authorization CVEs (gitea CVE-2025-68941, next.js CVE-2025-29927) that Semgrep OSS, CodeQL, and Cycode's own rule engine all missed. Confirmed findings that generalize are auto-promoted into deterministic SAST rules so a class caught once re-runs at near-zero cost.

> Cycode introduces Agentic Code Scanning, a fourth scanning dimension that uses agent orchestration to catch authorization and logic-level vulnerabilities rule engines structurally can't express, backed by per-scan tuning, attack-chain synthesis grounded in its Context Intelligence Graph, automated remediation workflows, and a self-improving rule flywheel.

└──▷ WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Agentic Workflows for automated remediation NEW

77

Agentic Workflows consume exploitability-qualified findings and automate remediation: teams define triggering events, agent action sequences, and confidence thresholds once, and the workflow fires automatically — routing findings to the owning engineer and returning fixes as reviewable pull requests with a full audit trail.

— Describes configurable triggers and thresholds, but no exact syntax.

Introducing Agentic Code Scanning: The H…

02 Agentic Code Scanning as a fourth scanning dimension NEW

73

An agent-orchestrated layer that plans from a threat model, ranks code by vulnerability class, and disproves every finding before surfacing it, running above the three existing deterministic/AI scanning tiers. This enables detection of logic-level flaws like broken access control that have no matchable pattern for rule engines. Across a ten-repository, six-language benchmark using an affordable open-weights reasoning model, it detected authorization CVEs gitea `CVE-2025-68941` and next.js `CVE-2025-29927` that Semgrep OSS, CodeQL, and Cycode's own rule engine all missed.

— Mechanism and named CVE benchmark, but no direct usage syntax given.

Introducing Agentic Code Scanning: The H… | Introducing Agentic Code Scanning: The H…

03 **Attack-chain synthesis via Context Intelligence Graph** NEW

65

Cycode links related findings across files into a single multi-step exploit path, carrying data flow from untrusted source to dangerous sink and tagging the resulting risk category (e.g. Code Injection — Critical). Attack chains are grounded in Cycode's Context Intelligence Graph, using real call relationships, ownership, and reachability rather than speculative model inference.

— Clear mechanism and named risk-category example, no direct action steps.

Introducing Agentic Code Scanning: The H…

04 **Per-scan mode, model, and scope configuration** NEW 60

Teams can configure three independent dials per scan — mode, model, and scope — to target expensive agentic reasoning only at code classes that rules cannot describe, while keeping deterministic scanning on every commit. Every scan records what ran, when, on which code, and with which model, for auditability.

— Names the three config dials but no exact command or file syntax.

Introducing Agentic Code Scanning: The H… | Introducing Agentic Code Scanning: The H…

THINNER COVERAGE BELOW

05 **Self-improving rule flywheel** NEW 58

Agentic Code Scanning findings that generalize are automatically converted into new deterministic SAST rules (Dimension 1/2), so a vulnerability class caught once by the model re-runs at near-zero cost indefinitely, compounding coverage as repository count grows.

— Explains the promotion mechanism but no configuration detail given.

Introducing Agentic Code Scanning: The H… | Introducing Agentic Code Scanning: The H…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Endor Labs

2 RELEASES · 2026-09-03

SAST / SCA / Code 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

Package Firewall intercepts every package install request and returns a 403 for malicious, vulnerable, or non-compliant packages before they touch disk, backed by a real-time malware feed scanning new npm, PyPI, and Go uploads within minutes and cooldown enforcement on fresh versions.

> Endor Labs launched Package Firewall, a new enforcement layer that blocks malicious, vulnerable, and non-compliant packages before they reach disk, backed by a real-time malware feed, policy engine integration, flexible deployment modes, cooldown enforcement, and full install logging.

└──▷ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 Artifact manager and endpoint deployment modes NEW 70

Package Firewall integrates with private artifact managers such as JFrog Artifactory by pointing a remote repository at the Package Firewall URL with no developer-machine changes, and supports an Endpoint Mode where configuration is pushed to developer machines via MDM and to CI runners via a config snippet, requiring no manual developer action.

– Names concrete deployment paths (Artifactory, MDM, CI snippet) as starting points

Introducing Package Firewall | Blog | En…

THINNER COVERAGE BELOW

02 Package Firewall blocks installs before disk write NEW 58

Package Firewall intercepts every package install request and returns a 403 for flagged packages — the package never touches disk — while clean packages are redirected transparently to the upstream registry.

– Explains mechanism but no config surface to act on directly

Introducing Package Firewall | Blog | En… launch-20260903-d7431612

03 Real-time malware feed with multi-pass analysis NEW 55

A real-time malware feed continuously scans every newly uploaded package across npm, PyPI, Go, and other ecosystems using multi-pass LLM, static, and behavioral

analysis, classifying packages within minutes of upload.

— Names ecosystems and analysis techniques but nothing a reader configures

Introducing Package Firewall | Blog | En… launch-20260903-d7431612

04 **Configurable cooldown periods for new package versions** NEW

50

Package Firewall enforces cooldown periods, blocking newly published package versions for a configurable window regardless of requester, addressable at scale via policy.

— Describes behavior and policy control but no exact config key

Introducing Package Firewall | Blog | En…

05 **Install event logging in UI and API** NEW

50

Every install event — package name, version, time, and reason — is logged and visible in the Endor Labs UI and via the API.

— Names logged fields and access surfaces but no endpoint path

Introducing Package Firewall | Blog | En…

06 **Policy engine enforcement for CVEs and license compliance** NEW

45

Package Firewall enforces the full Endor Labs policy engine at install time, blocking packages with known CVEs or that violate license compliance policies, beyond malware detection alone.

— States policy scope but no named policy fields or endpoints

Introducing Package Firewall | Blog | En… launch-20260903-d7431612

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Horizon3 NodeZero

2 RELEASES · 2026-08-31 → 2026-09-01

Exploitation & C2

Horizon3 NodeZero is an autonomous penetration testing platform that continuously finds and exploits real attack paths in your environment to identify and prioritize critical vulnerabilities.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Adds autonomous WebApp and Azure Cloud pentest op types—crawling and testing authenticated production/staging web apps and assessing Azure control planes plus VM workloads in a single assumed-breach run with proof-based findings, extending autonomous exploitation into two attack surfaces it could not reach before.

> NodeZero expanded its assessment surface with new WebApp and Azure Cloud pentest op types, an Azure Entra ID AD Password Audit, Australian data residency, and enhanced Active Directory/Identity and op-tracking capabilities.

└──▷ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 Azure Cloud pentest op type NEW

80

New Azure Cloud Pentest op type assesses Azure control planes and VM workloads in a single assumed-breach assessment, combining Azure API enumeration with an Internal pentest. It lets practitioners set up an Azure connection, run cloud pentests, manage connections, and perform manual cleanup, launched via Launch Op > Azure Cloud Pentest with Azure credentials/scope and a combined Azure API + Internal pentest assessment mode.

Assess an Azure environment end-to-end — from control plane enumeration to lateral movement on VMs — in a single assumed-breach op.

📍 In NodeZero, go to Launch Op › Azure Cloud Pentest, supply your Azure credentials/scope, and select the combined Azure API + Internal pentest assessment mode.

— Combines mechanism, scope and UI path across releases `snapshot-20260831` `snapshot-20260901`

02 WebApp pentest op type NEW

70

New WebApp pentest op type crawls, discovers, and tests production, staging, and private web applications with authenticated workflows, automated route discovery, and proof-based findings. It is launched via Launch Op > WebApp Pentest by entering a target URL and configuring authenticated workflow credentials.

Run a WebApp pentest against a staging application to discover routes and get proof-based findings before a production release.

> 📍 In NodeZero, go to Launch Op › WebApp Pentest, enter your target URL, configure authenticated workflow credentials, and start the assessment.

— UI path and workflow named, no API detail `snapshot-20260901`

---

03 **AD Password Audit for Azure Entra ID** NEW — 65

New AD Password Audit feature audits Active Directory passwords within Azure Entra ID to identify weak or compromised credentials, accessed via Azure Entra ID > AD Password Audit by connecting an Azure tenant and launching the audit.

Audit Active Directory passwords in your Azure Entra ID environment to identify weak or compromised credentials.

> 📍 In the NodeZero console, navigate to Azure Entra ID › AD Password Audit, connect your Azure tenant, and launch the audit.

— Clear navigation path but no scoring or output detail `snapshot-20260831`

THINNER COVERAGE BELOW

04 **Activity and Artifacts tabs in op views** NEW — 45

The Action Log now extends from Real-Time View into post-op results, and both the Real-Time and post-op views gain a dedicated Artifacts tab to deconflict NodeZero activity from malicious attacks.

— Names the tabs but not how to reach them `snapshot-20260901`

---

05 **Active Directory and Identity assessment enhancements** IMPROVED — 35

AD & Identity assessments gain a new dedicated Identity Details view, Kerberos ticket extraction during credential dumping, and safer authentication workflows.

— Named additions but no mechanism or steps given `snapshot-20260901`

---

06 **Australia data residency** NEW — 30

Platform traffic and pentest data can now be kept within the Australian region for data-sovereignty compliance.

— States the capability without setup steps `snapshot-20260901`

└──▷ ALSO FROM THESE RELEASES

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 adscan

1 RELEASE · 2026-09-01

Exploitation & C2

Free Active Directory pentesting tool for Linux. Automates AD and LDAP enumeration, Kerberoasting, ASREPRoast, password spraying, ADCS/ESC1, DCSync, RBCD, gMSA, shadow credentials, NTLM relay and credential dumping across 104 techniques, mapping BloodHound-compatible attack paths to Domain Admin.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Attack paths now cross domain and forest boundaries as a unified graph—cross-forest Kerberoasting, AS-REP Roasting, and ADCS ESC8/ESC11 relay routes across trusts—plus automated cross-forest escalation and a minimal-footprint krbtgt-only DCSync mode for proving full domain compromise.

> ADscan v11.3.0 adds a built-in didactic/teaching mode for techniques, seasonal password spraying, a minimal-footprint krbtgt-only DCSync mode, cross-forest attack path analysis with automated escalation routes, LDAP-attribute credential harvesting as a validated attack step, and independently-verifiable evidence blocks in reports.

└──▷ WHAT SHIPPED · 7 FEATURES most completely described first ? what's the number?

01 Didactic mode for attack techniques NEW

91

Adds `set explain_level off|basic|deep` to control teaching verbosity, and `explain <technique>` (e.g. `explain kerberoasting`, `explain "ADCS ESC1"`) to display a teaching card — MITRE ATT&CK ID, the equivalent manual command, and detection Event IDs — for any technique without actually executing it.

Before a Kerberoasting step in a lab, display its full teaching card — MITRE ID, manual equivalent, and detection Event IDs — without triggering the attack.

```
$ explain kerberoasting
```

— Names exact commands, flags, and a runnable teaching example. v11.3.0

02 Seasonal password spraying strategy NEW

80

Adds a `month_season` spray strategy to the scan config for seasonal password spraying (e.g. March2026, Verano2026), with per-domain language inference (English/Spanish) and last-password-change awareness when domain credentials are available; runs automatically in `adscan ci`.

Run a fully unattended CI scan that includes both month and season spray passes against a domain, using the month_season strategy.

```
$ adscan ci auth --type audit --interface eth0 --domain
corp.local --dc-ip 10.0.0.1 -u alice -p 'S3cr3t!'
```

— Names the config strategy and gives a runnable CI command. v11.3.0

03 **Independent verification block in findings** NEW 78

Reports and the web platform now include a 'Verify this finding independently' block per attack step with native Windows/PowerShell (`Get-ADUser` , `Get-Acl` , `dsacls` , `certutil`) and Linux equivalents pre-filled with the real DC IP and domain name, covering Kerberoasting, AS-REP Roasting, DCSync, ADCS ESC1, and the GenericAll/GenericWrite/WriteDACL/WriteOwner ACL family.

— Names concrete verification commands pre-filled per finding. v11.3.0

04 **Cross-forest attack path analysis** NEW 65

Attack paths now cross domain and forest boundaries in multi-domain workspaces, analysed as a unified graph, including cross-forest Kerberoasting, AS-REP Roasting, and ADCS ESC8/ESC11 relay routes across trusts.

— Names techniques and scope but no command to run. v11.3.0

05 **Automated cross-forest escalation routes** NEW 65

Automated cross-forest escalation now ships two routes: from a compromised trusted forest into the trusting forest via cross-organization TGT-delegation trust, and from a compromised child domain up to its forest root (RaiseChild), each triggering post-compromise on the reached domain.

— Names the two escalation mechanisms with no runnable example. v11.3.0

THINNER COVERAGE BELOW

06 **Credential harvesting from LDAP description/info attributes** NEW 55

Plaintext credentials in account `description` or `info` LDAP attributes are now a validated, end-to-end attack step — ADscan reads the attribute, recovers the credential, verifies login, and traces the route across forest boundaries.

— Names the LDAP fields but gives no command or flag. v11.3.0

07 **krbtgt-only DCSync replication mode** NEW 50

DCSync now supports a 'krbtgt only' replication mode, capturing the golden-ticket key without pulling every account — a minimal-footprint way to prove full domain compromise.

— Describes the mechanism but no command or flag given. v11.3.0

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 Binary Ninja

1 RELEASE · 2026-09-03

Reverse Engineering

Binary Ninja is a modern reverse engineering platform with a scriptable and extensible decompiler.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

Binary Ninja 6.0 ships a first-party MCP server and a Binary Similarity diffing platform (Google BinDiff and WARP providers) with automatic annotation transfer between related binaries, plus TMS320C6x DSP support and up to 2.16x faster analysis at 55% less peak memory.

> Binary Ninja 6.0 (Krypton) adds a first-party MCP server for LLM client integration, a new Binary Similarity diffing platform for Ultimate, TMS320C6x architecture support, and up to 2.16× faster analysis with lower memory use.

└──▷ WHAT SHIPPED · 12 FEATURES most completely described first ? what's the number?

01 MCP server for LLM client integration NEW

95

Binary Ninja ships a built-in GUI MCP server, enabled via the `ui.mcp.enabled` setting (restart required) or `Plugins > MCP > Start Server`, listening by default at `http://127.0.0.1:24642/mcp`. Access can be locked down with the `ui.mcp.token` setting requiring clients to send a matching bearer token, and `Plugins > MCP > Copy Connection Info` copies the exact URL and authorization header for the active session. A standalone `binaryninja_mcp` headless stdio-based MCP server binary is also shipped for Commercial and Ultimate on macOS and Linux.

Enable the MCP server on every launch so an LLM client (e.g. Claude Desktop) can always connect to your Binary Ninja session.

> 📍 In Binary Ninja, go to Settings, search for `ui.mcp.enabled`, and set it to true, then restart. Navigate to Plugins › MCP › Copy Connection Info to copy the URL and bearer token for your client config.

Restrict MCP access to authorised clients only by requiring a bearer token.

> 📍 In Binary Ninja Settings, search for `ui.mcp.token` and set it to a strong secret string. Clients must then include a matching `Authorization: Bearer <token>` header when connecting to `http://127.0.0.1:24642/mcp`.

— Exact settings, menu paths, URL and headless binary name all given

Binary Ninja 6.0 (Krypton) (2026-09-03)

02 **Binary Similarity diffing platform** NEW 75

A new Binary Similarity feature (Ultimate only) compares related binaries using pluggable providers, shipping with Google BinDiff and WARP providers, and automatically transfers annotations between binaries. It is exposed through Python, C++, and Rust APIs, enabling headless similarity sessions and registration of custom providers and resolvers.

— Providers and APIs named but no exact call syntax given

Binary Ninja 6.0 (Krypton) (2026-09-03)

THINNER COVERAGE BELOW

03 **Refactored calling convention system** IMPROVED 55

The calling convention system was refactored to properly represent structure parameters and return values, with new Go and Pascal calling conventions added.

— Names new conventions and behavior change, no exact API

Binary Ninja 6.0 (Krypton) (2026-09-03)

04 **BASE sampling mode for faster base address detection** IMPROVED 50

Introduces a new BASE 'sampling' mode, enabled by default, for faster automated base address detection especially on large files; the previous IL Analysis Mode remains available when higher accuracy is needed.

— Names both modes and default but no toggle command

Binary Ninja 6.0 (Krypton) (2026-09-03)

05 **Faster analysis with lower memory use** IMPROVED 50

Analysis runs up to 2.16× faster and uses up to 55% less peak memory versus version 5.3 across representative binaries.

— Concrete benchmark numbers but no reproducible steps

Binary Ninja 6.0 (Krypton) (2026-09-03)

06 **Multiple global pointer registers per architecture** IMPROVED 45

Binary Ninja now supports multiple independent global pointer registers per architecture, such as TriCore's four global data region pointers, improving decompilation accuracy for such targets.

— Explains mechanism and example target but no steps

Binary Ninja 6.0 (Krypton) (2026-09-03)

07 **HLIL_STRUCT_INITIALIZE IL operation** NEW 45

Adds `HLIL_STRUCT_INITIALIZE` IL operation for representing structure initializers in HLIL.

— Named IL operation but usage context not described

Binary Ninja 6.0 (Krypton) (2026-09-03)

---

08 **Free edition gains AArch64, MCP, and Linux ARM64 build** NEW 40

The free edition adds AArch64 (ARMv8) architecture support and the MCP server, plus a new Linux ARM64 package.

— Names the three additions but no configuration detail

Binary Ninja 6.0 (Krypton) (2026-09-03)

---

09 **TMS320C6x architecture support** NEW 35

Adds TMS320C6x (DSP) architecture support, bringing Binary Ninja Ultimate to 19 first-party architectures.

— Names the architecture but no usage detail Binary Ninja 6.0 (Krypton) (2026-09-03)

---

10 **Extension Manager replaces Plugin Manager** IMPROVED 35

Replaces the Plugin Manager with a new Extension Manager UI.

— Names the new UI but no navigation or capability detail

Binary Ninja 6.0 (Krypton) (2026-09-03)

---

11 **Python 3.13 bundled runtime** IMPROVED 30

Upgrades the bundled Python to 3.13, now also included on Linux.

— Version and platform named, no further detail Binary Ninja 6.0 (Krypton) (2026-09-03)

---

12 **New user onboarding wizard** NEW 25

Adds a new user wizard to ease migration for new and upgrading users.

— Appears on launch but mechanism unspecified Binary Ninja 6.0 (Krypton) (2026-09-03)

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · September 4, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe