THE CYBER TOOLCHAIN NO. 007

Head

THE WEEK'S TOP PICKS

PUBLISHED AUGUST 14, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

EDITIONS tail | grep | **head** | diff | uniq

The best tooling updates that shipped this week.

// HOW THESE PICKS ARE MADE

Every feature release from the 874 tools on our watchlist goes into the daily newsletter. Once a week we read the whole field side by side and pick the few updates worth your attention, 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 PICKS

01	**adscan**	Exploitation & C2
02	**Elastic Security Labs**	SIEM
03	**Radare2**	Reverse Engineering
04	**Wiz Code**	SAST / SCA / Code Security
05	**Tenzir**	Detection Engineering

01 adscan

4 RELEASES · 2026-07-16 → 2026-08-10

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

Adds cross-forest privilege-escalation chains, surfacing cross-forest Kerberos delegation and MSSQL linked-server compromise across trust boundaries as a single blast radius, plus Tier 0 membership counts and ADSCAN_TRUST_SCOPE trust-collection control. The prior release also executes AD CS ESC5 end to end and generates MITRE-mapped coverage and reports.

> Across four releases, ADscan shipped an integrated password-cracking pipeline with a bundled 94-million-entry wordlist, full end-to-end ADCS ESC5 exploitation, SQL Server and Shadow Credentials attack coverage, cross-forest trust escalation modeling, and resumable scans with native macOS/ARM64 support, alongside new `harvest`, `jobs`, `generate_report`, `writeup`, and `adscan execute attack_paths` commands. The window also carried several breaking changes: bundled external tools were removed in favor of a native engine, nine REPL commands were renamed, and the tool now refuses to run outside its container.

└──▷ WHAT SHIPPED · 38 FEATURES most completely described first ? what's the number?

01 Cross-forest trust scope and escalation modeling NEW

88

`ADSCAN_TRUST_SCOPE` (values `all` or `origin`) controls which domains are collected across forest trusts, overridable per scan via the web scan-setup domain-enumeration policy. Cross-forest Kerberos ticket delegation is now surfaced as an escalation route, drawing both forests as one blast radius with remediation to disable the delegation. MSSQL linked-server-chain compromise now extends across forest boundaries, coupling SYSTEM-on-DC to directory replication in the target trusted domain rather than stopping at the controller. Trust relationships are always mapped regardless of collection scope, and domain collection follows engagement type: CTF scans collect all reachable domains including newly unlocked trusted forests unattended, while audit scans default to the origin domain.

Run an audit scan locked to the origin domain only, preventing collection of trusted forests outside the agreed legal scope.

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

— Names env var, values and per-scan override with a runnable example. v11.1.0

02 **Integrated password-cracking pipeline** NEW 86

Adds password-strength auditing built into the scan: credentials captured via broadcast poisoning, Kerberos service tickets, and AS-REP hashes are cracked in the background using the bundled 94-million-entry 'ADscan AD Audit Wordlist' and a client-targeted wordlist generator, with recovered passwords fed back into the attack path automatically. A four-rung selectable cracking-effort ladder is calibrated to host hardware, with accurate pre-commit time estimates and live progress/ETA during cracking; cracking-effort and poisoning controls are exposed in the YAML scan-config file and the web platform.

Run a fully unattended CI scan with cracking effort and poisoning capture configured ahead of time via the YAML scan-config file.

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

— Names wordlist size, ladder and config surfaces; exact YAML key missing. v10.0.0

03 **`adscan execute attack_paths` subcommand** NEW 81

New `adscan execute attack_paths` subcommand runs a single saved attack path or step against a workspace without a full rescan; it resolves credentials automatically for already-owned accounts, and a companion re-arm action resets step status to re-run a path.

Re-run a single saved attack path against a workspace after remediating one step, without repeating the full scan.

```
$ adscan execute attack_paths
```

— Runnable command with clear mechanism and example. v10.1.0

04 **REPL flag-form arguments and tab-completion** IMPROVED 78

REPL commands now accept flag-form arguments `-d` / `--dc-ip`, `-u`, `-p`, `--host`, `--service` alongside the existing positional form, with tab-completion for both flag names and values.

Use flag-form arguments in the REPL for scripted or tab-completed input instead of positional syntax.

```
$ start_auth -d 10.0.0.1 -u alice -p 'S3cr3t!' --host corp.local
```

— Exact flags named with a runnable example. v10.1.0

---

05 **`harvest` command for credential inventory** NEW 73

New `harvest` command provides a single view of every credential obtained during a scan (captured, cracked, and sprayed), each tagged with its privilege tier and reachable assets, plus an end-of-scan harvest summary.

Pull up every credential recovered so far — captured hashes, cracked passwords, and sprayed accounts — with privilege tier and reach, mid-engagement.

```
$ harvest
```

— Named command with runnable example and clear scope. v10.0.0

---

06 **Client-logo branding for reports** NEW 70

New `adscan deliver --client-logo` and `adscan ci --client-logo` flags accept PNG, SVG, or JPG to embed a client logo in both the free and PRO report covers.

Generate a client-branded exposure report during a non-interactive CI scan, embedding the client's logo in both the HTML and PDF output.

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

— Exact flags and accepted formats named, runnable example given. v11.0.0

---

07 **Telemetry opt-out controls** NEW 70

Adds `--no-telemetry` and `--offline` flags to `adscan start`; also restores `set telemetry off` as a per-workspace setting and introduces `set telemetry off global` for an operator-wide default.

— Exact flags and REPL commands named. v11.0.0

---

08 **SQL Server attack coverage and automation** NEW 69

Adds SQL Server attack coverage: linked-server lateral movement across domain and forest trusts, automated command execution with escalation to SYSTEM, and arbitrary file read via bulk import, each surfaced as a validated attack step in the report. The SQL Server post-authentication workflow now runs automatically in unattended scans instead of defaulting to skip.

— Names three attack techniques and an automation change, no flag given. v10.1.0

09 **`writeup` command for lab/CTF documentation** NEW 68

New `writeup` command captures port data, directory contents, the full attack chain as an editable Mermaid diagram, timestamped steps with credential provenance, flags, and dead ends into a lab/CTF writeup; it runs automatically on lab and CTF scans.

— Named command with detailed contents, no example given. v11.0.0

10 **Escalation account handling for MSSQL and HasSession** IMPROVED 67

MSSQL SYSTEM-escalation follow-up and HasSession exploitation now offer a choice between minting a fresh privileged account or elevating one already under the operator's control. HasSession escalation elevates to the host's own local Administrators group on member servers instead of always adding to Domain Admins. At workspace exit, ADscan prompts whether to delete any account it created during an MSSQL follow-up (default: yes); declining records the exact removal command so nothing is silently left behind.

— Full mechanism and default behaviour described, no flag named. v11.1.0

11 **End-to-end ADCS ESC5 exploitation** NEW 67

ADscan now executes AD CS ESC5 end to end: it backs up the CA private key, forges a certificate offline, and authenticates with it from the CA host's machine account or a recovered local-admin password.

— Full attack chain mechanism described, no CLI surface named. v11.0.0

12 **`jobs` command for background task control** NEW 65

New `jobs` command lists, reviews, and stops long-running background tasks such as broadcast-name-resolution poisoning capture and password cracking.

Review and stop background cracking or poisoning jobs without interrupting the active scan.



— Named command with example and scope of tasks listed. v10.0.0

13 **Broadcast name-resolution poisoning capture** NEW 65

Adds broadcast name-resolution poisoning capture (LLMNR, NBT-NS, mDNS) as a background task, surfaced as a finding and linked attack step, with an explicit consent panel in interactive mode and automatic start in unattended scans.

— Names protocols and consent behaviour, no toggle key given. v10.0.0

14 **`generate_report` command for exposure reports** NEW 62

New `generate_report` command produces a free self-contained HTML report and a paginated A4 PDF exposure report from any scan; audit-mode scans now run it automatically.

— Named command and output formats, but no example shown. v11.0.0

15 **Built-in engine replaces external command-line tools** BREAKING 62

Password spraying, share enumeration, SMB/LDAP hardening checks, local credential verification, and account cleanup now run on ADscan's built-in engine with consistent proxy support, removing the dependency on bundled external command-line tools. Bundled external third-party command-line cyber tools have been removed; any workflow that relied on them directly must now use ADscan's native engine equivalents.

— Lists five affected functions and the removal, no migration command. v10.0.0

16 **Collapsed attack paths for interchangeable accounts** NEW 61

Attack paths with interchangeable intermediate accounts now collapse into a single finding that names and counts those accounts, reducing report duplication. The operator can choose which interchangeable accounts to target — several at once for a roast or spray, or one for a direct step — with unattended runs picking the strongest candidate automatically.

— Clear before/after and operator choice, no flag named. v11.0.0

17 **Writable-share credential bait** NEW 61

Adds opt-in writable-share credential bait: plants a bait file on writable shares and captures resulting authentication attempts as a background job that does not block the scan; configurable via new scan-configuration toggles in the CLI and web console.

— Mechanism and opt-in toggles named, no exact key given. v10.1.0

18 **COVERAGE.md technique mapping** NEW 60

New `COVERAGE.md` lists all 103 Active Directory techniques implemented, each mapped to MITRE ATT&CK and marked executed, detected, or declined for safety; it is generated from the product catalog so it cannot drift.

— Names artifact and generation source, no command to run it. v11.0.0

19 **Remediation guidance limited to native Microsoft tooling**

IMPROVED 60

Remediation guidance is now written entirely in native Microsoft tooling (`dsacls` , RSAT Active Directory module, `Get-*` / `Set-*` cmdlets) with no third-party product names in client-facing report text.

— Names exact tools used, but no example command. v10.1.0

THINNER COVERAGE BELOW

20 **Native Certificate Services attack implementation** IMPROVED 58

Certificate Services attacks now run through a native implementation, including detection of delegated CA-management rights and detection of Extended Protection on the web-enrollment endpoint, which suppresses false ESC8 findings when enrollment is hardened.

— Names two detections and a false-positive fix, no CLI surface. v10.1.0

21 **Interactive SMB host-enrichment scope prompt** NEW 55

On large directories, ADscan now prompts the operator to choose the SMB host-enrichment scope: the recommended host set, a custom count, all hosts, or skipping enrichment entirely. Reports and the platform state how many hosts were enriched out of the reachable total.

— Describes the prompt options but names no CLI flag. v11.1.0

22 **Engagement-aware Force Change Password automation** IMPROVED 55

Force Change Password steps now follow engagement type in unattended runs: CTF scans reset the target password and continue automatically, while audit scans skip the reset unless an operator explicitly approves it. Machine account password resets are refused outright in every mode.

— Clear before/after behaviour, no direct flag named. v11.1.0

23 **Native macOS and Apple Silicon support** NEW 53

Adds native macOS and Apple Silicon support; LITE and PRO images are now published for ARM64/Apple Silicon in addition to x86-64.

— Names exact image variants and architectures supported. v10.0.0

24 **Resumable scans** NEW 52

Adds resumable scans — interrupted scans (crash, disconnect, or deliberate stop) can resume from the last completed phase or host; the web platform can stop and resume a running scan gracefully.

— Describes resume points, no exact command shown. v10.0.0

25 **Credential capture through a routed pivot** NEW 49

Adds credential capture through a pivot: NTLM authentication can be captured across a routed tunnel, with per-target selection of the callback address.

— Names mechanism and per-target config, no exact flag. v10.1.0

26 **Cross-scan exposure comparison view (Enterprise)** NEW 47

The Enterprise platform adds a cross-scan exposure comparison view that plots exposure across a workspace's scans and lists findings and attack paths that appeared, closed, or carried over.

— Describes the view but no navigation path given. v11.0.0

27 **Shadow Credentials escalation** NEW 47

Adds Shadow Credentials escalation: detects when it can write a key credential to a target account and executes the full authentication chain end to end.

— Describes mechanism briefly, no config surface named. v10.1.0

28 **Report status and exposure scoring refinements** IMPROVED 45

Attack paths now carry more precise result states: `'Partially Validated'`, `'Not Executed for Safety'`, and `'Attack Surface Reduced — Hardening Observed'`. The AD exposure score now combines proof strength with per-technique exploitability.

— Names the three result states and scoring inputs, no formula given. v10.1.0

29 **ConvertFrom-SecureString secret recovery** NEW 44

Adds recovery of secrets protected with PowerShell `ConvertFrom-SecureString -Key`, chained into share-password findings.

— Names exact PowerShell cmdlet but thin mechanism. v10.1.0

30 **Privilege Blast Radius view** NEW 44

Adds Privilege Blast Radius: reports, the web platform, and the CLI now show the full set of assets and capabilities a single compromised account reaches.

— Clear description, no navigation or command given. v10.0.0

31 **Extended-rights escalation modeling** NEW 43

Extended-rights escalation (forced password reset, DCSync, LAPS password read) is now modeled per target and chained into attack paths.

— Names three rights but limited mechanism detail. v10.1.0

32 **HTTP SPN accounts flagged as relay surface** NEW 40

Service accounts carrying an HTTP SPN are now flagged as a Kerberos web-application relay and coercion surface, with native remediation guidance included in the report and platform.

— Names the technique but no config surface. v11.0.0

33 **GPO-to-code-execution chaining** NEW 38

GPO control is now chained through to code execution via a scheduled task on machines the GPO applies to.

— Brief mechanism description, no config surface. v10.1.0

34 **Tier 0 membership counts in reports** NEW 35

Reports and the platform now show how many accounts already hold Tier 0 privilege via group membership alongside the count reachable via attack paths, with already-privileged accounts excluded from the attack-path figure.

— Clear behaviour but no surface to act on directly. v11.1.0

35 **REPL commands renamed to native names** BREAKING 32

Nine REPL commands were renamed to native names; the previous command names are no longer accepted.

— States the break but doesn't list the nine renamed commands. v10.1.0

36 **Web ports 8080/8443 added to service sweep** IMPROVED 30

Web ports 8080 and 8443 are now included in the important-ports sweep and folded into service discovery.

— Names exact ports, minimal further detail. v10.1.0

37 **Container-only execution enforced** BREAKING 30

Running the ADscan core directly on the host machine is now refused; the tool runs only inside its container.

— States the restriction but gives no migration guidance. v10.1.0

38 **'Blocked by Active Controls' category removed** BREAKING

The 'Blocked by Active Controls' report category was removed.

— Bare removal notice with no replacement described. v10.1.0

└──▷ BREAKING ON UPGRADE

- ! Nine REPL commands were renamed to native names; the previous command names are no longer accepted.
- ! The 'Blocked by Active Controls' report category was removed.
- ! Running the ADscan core directly on the host machine is now refused; the tool runs only inside its container.
- ! Bundled external third-party command-line cyber tools have been removed; any workflow that relied on them directly must now use ADscan's native engine equivalents.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Elastic Security Labs

2 RELEASES · 2026-08-07 → 2026-08-11

SIEM

Elastic Security Labs is a research division that publishes threat intelligence, malware analysis, and security research to help defenders detect and respond to cyber threats.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

Ships a dependency-free bash/PowerShell hook collector that streams every Cursor AI coding-agent tool call to Elasticsearch as ECS-conformant JSONL, letting teams run ES|QL queries to find, for example, which hosts ran an agent that read a .pem file. A companion CEL integration snapshots .npmrc to catch silent removal of npm's min-release-age cooldown.

> Elastic Security Labs shipped an end-to-end auditing pipeline for Cursor AI coding agent activity — a hook collector script, an MDM-less distribution mechanism, and an Elastic Agent integration for ES|QL queries — plus a CEL-based detector that catches silent removal of npm's min-release-age cooldown setting.

└──▷ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Bash/PowerShell hook collector for Cursor agent tool calls NEW

95

A 280-line dependency-free bash script `log-tool-calls.sh` (plus a PowerShell port for Windows), published in the `elastic/elasticsearch-labs` repository, records every Cursor AI agent tool call as JSONL to a date-rotated log file readable only by the owner (`chmod 0600`). Each entry promotes top-level fields `hook_event_name` , `tool_name` , `command` , `file_path` , `mcp_server` , `model` , `session_id` , `duration` , `duration_ms` plus an ECS-conformant `event` object (`event.category` , `event.type` , `event.outcome`), and is tagged with a detected `agent` / `ide` surface (`cursor` , `cursor-cli` , `remote`) by inspecting `CURSOR_VERSION` , `CURSOR_CODE_REMOTE` , `VSCODE_PID` , `VSCODE_CWD` , `VSCODE_IPC_HOOK` . The same `hooks.json` also covers the headless `cursor-agent` CLI used in CI, and local logs are retained for 30 days with opportunistic deletion, with Elasticsearch as the system of record.

Deploy the hook collector system-wide on macOS so both the Cursor IDE and `cursor-agent` CLI log every tool call without MDM.

```
$ sudo mkdir -p /usr/local/share/ai-hooks
sudo cp log-tool-calls.sh /usr/local/share/ai-hooks/log-tool-
calls.sh
sudo chmod +x /usr/local/share/ai-hooks/log-tool-calls.sh
sudo tee "/Library/Application Support/Cursor/hooks.json"
<<'EOF'
{
  "version": 1,
  "hooks": {
    "sessionStart":          [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "beforeShellExecution":  [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "afterShellExecution":   [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "beforeMCPExecution":    [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "afterMCPExecution":     [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "postToolUse":           [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "postToolUseFailure":    [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "afterFileEdit":         [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "beforeReadFile":        [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "subagentStart":         [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "subagentStop":          [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "sessionEnd":            [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "stop":                  [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }]
  }
}
EOF
```

Verify the CLI surface is captured separately from the IDE after deployment by running a headless agent call and checking the log.

```
$ cursor-agent --print "say hello"
tail -1 ~/.config/ai-hooks/logs/tool-calls-$(date -u +%F).jsonl
```

— Full script, config file and verification commands given

13 million tool calls: auditing every AI…

02 CEL integration detects silent npm cooldown removal NEW

85

A ~40-line CEL integration in Elastic Agent snapshots every `.npmrc` on a 6-hour heartbeat and sets `cooldown.absent = true` in the ingest pipeline when `min-release-age` disappears between snapshots. It uses the file() and dir() CEL functions to enumerate user home directories (`/Users` , `/home`) and machine-global npm config paths (`/opt/homebrew/etc/npmrc` , `/usr/local/etc/npmrc` , `/etc/npmrc` , `/usr/lib/node_modules/npm/.npmrc`) across macOS and Linux, sets `resource.url: file:///dev/null` to satisfy CEL's required URL field, and takes global paths and poll interval via the "Custom request cursor" YAML field in Fleet UI. Registry auth token lines (`//registry.npmjs.org/:_authToken=...`) are filtered in the ingest pipeline before indexing, and per-file SHA-256 hashes (`file.hash.sha256`) are emitted so in-place rewrites that strip `min-release-age` are caught even without a path change. The design note explains that the Custom Logs Filestream integration was unsuitable because files under 1024 bytes are held back under the default fingerprint identity strategy, `filestream` only emits on append, and `clean_inactive` requires `ignore_older` — a poor fit for static config files that may sit unchanged for months.

— Exact field names and paths given, but no runnable setup command

The security signal log tailing can't se…

THINNER COVERAGE BELOW

03 MDM-less deployment via Cursor Cloud Distribution console NEW

55

Delivers the collector to devices outside MDM via Cursor's Cloud Distribution console, which pushes a self-bootstrapping hook command carrying the collector script gzipped and base64-encoded, verified against a pinned SHA-256 hash on every invocation.

— Mechanism named but no runnable command given 13 million tool calls: auditing every AI…

04 Elastic Agent filestream integration for agent log queries NEW

45

An Elastic Agent filestream integration configuration parses the JSONL hook logs into structured fields, enabling ES|QL queries across all agent activity — for example, identifying which hosts ran an agent that read a `.pem` file in a given time window.

— Names ES|QL and use case but no config or query shown

13 million tool calls: auditing every AI…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Radare2

1 RELEASE · 2026-08-07

Reverse Engineering

UNIX-like reverse engineering framework and command-line toolset

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

Adds a Rust v0 symbol demangler that makes modern Rust binaries readable, plus new binary-resource extraction commands (iU/ix, rabin2 -xU) and struct synthesis from pointer-argument access patterns (afts), with PPC64 ELFv2 calling-convention and WebAssembly type support.

> Radare2 6.2.0 adds a generic binary resource extraction system, an initial Rust v0 demangler, struct type synthesis from pointer-argument patterns, and broader calling-convention analysis for PPC64 and MIPS, alongside several breaking changes including a Capstone v4 minimum and stricter r2agent authentication.

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

01 Generic binary resource extraction system NEW

95

Adds the `iuo` command to open binary resources inline, `ix` to extract sections, segments, and resources, and `iU` (with `iUx` extraction) plus the `rabin2 -xU` flag for listing and extracting resources; the old `iR` command is renamed to `iU` and the old `iU` is moved to `iu`. Backed by a new generic resource API supporting GLib Resources in ELF, Macho code signatures, PE/CLR resources, PE resources from MDMP, OS/2 LE/LX resources, QNX LMF resources, and Pebble resources.

Extract all embedded resources from a PE binary directly from the command line.

```
$ rabin2 -xU malware.exe
```

— Names every command, flag and supported resource format with a runnable example 6.2.0

02 Struct type synthesis and field type propagation NEW

85

Adds struct type synthesis from pointer-argument access patterns via the `afts` command, and a `types.fields` config key to gate propagation of callee types into struct and union fields during `aaft`.

Synthesize struct types from pointer-argument access patterns after analyzing a function to recover unknown data structures.

```
$ r2 /bin/target
> aaa
> s sym.process_packet
> afts
```

— Runnable command example plus a named config key controlling the behaviour `6.2.0`

---

03 **Calling convention analysis across architectures** NEW

80

Adds automatic ppc64 ELFv2 calling convention selection from the ELF ABI, automatic MIPS calling convention selection from the o32/n32 ABI, and callconv definitions (clobber and preserve registers) for all supported architectures including PPC and ARM. Adds PPC64 TOC pointer base-register tracking gated by `RArchOpMask.STATEFUL` and wired into `aar`, plus new `r_anal_cc_argslot` and `r_anal_cc_argval` APIs for resolving calling-convention stack argument slots, `r_anal_cc_isclobber` for querying register clobbering, and reguse analysis hints with dyncc register role APIs.

— Names many specific APIs and architectures but is developer-facing, not a direct command

`6.2.0`

---

04 **New configuration keys for analysis and UI** NEW

68

Adds `anal.gpseed` to seed a PIC base register during ESIL analysis, `asm.cmt.reguse` to emit register-use comments in disassembly, `dir.binsdb` (from binpath/dirprefix) for PE/MDMP ordinal lookups, and `scr.drag` to enable click-and-drag scrolling in disasm, hex, and graph views.

Seed a PIC base register before ESIL-based analysis so position-independent code resolves correctly.

```
$ r2 /bin/firmware.elf
> e anal.gpseed=0x10000
> aaa
```

— Names four config keys with usage but limited mechanism detail `6.2.0`

05 **Stricter r2agent web session authentication** BREAKING

60

Adds auth protection to `r2agent` web session spawns via an HTTP auth file; empty lines in the auth file now block all access rather than allowing any password through, a behavior change from prior releases.

— Describes the mechanism and the exact behavior change, but no config path given 6.2.0

06 **Plugin API extensions** NEW

60

Adds the `Lcc` command to list all plugins registered with the new command API, `RArchOpMask.STATEFUL` and `RArchSession.reset` for stateful architecture decoding in plugins, and makes RBin demanglers pluggable and selectable.

— Names three distinct plugin-facing APIs with no usage example 6.2.0

THINNER COVERAGE BELOW

07 **Shell command chaining with && and ||** NEW

45

Adds `&&` and `||` short-circuit command chaining support in the r2 shell, including support for chaining after literal commands in the shell pipeline.

— Names the operators and where they apply but no worked example 6.2.0

08 **WebAssembly v128 and reference type support** NEW

40

Adds `v128`, `ref`, and `reftype` support for WebAssembly binaries.

— Names the specific Wasm types added but no usage shown 6.2.0

09 **New build system platform targets** NEW

40

Adds MiKNix platform support in the build system and a `dist/jslinux` target using temu/rv64 for JavaScript-based Linux builds.

— Names two build targets with minimal explanation of use 6.2.0

10 **Capstone v4 now minimum required version** BREAKING

35

Removes Capstone v3 support; Capstone v4 is now the minimum required version.

— States the version cutoff but no migration detail 6.2.0

11 **Rust v0 symbol demangler** NEW

30

Adds an initial Rust v0 symbol demangler, making Rust binaries easier to analyze.

— Only a bare description, no mechanism or usage given 6.2.0

12 **RBinPluginAny renamed and delinked from LibMagic**

BREAKING

30

`RBinPluginAny` is renamed to `"null"` and no longer uses LibMagic.

— Names the rename but gives no reason or migration path `6.2.0`

└──▷ BREAKING ON UPGRADE

! The `iR` command is renamed to `iU`; the old `iU` command is moved to `iu`; the `rabin2 -u` flag behavior is moved to an environment variable.

! Capstone v3 support is removed; Capstone v4 is now the minimum supported version.

! `RBinPluginAny` is renamed to `"null"` and no longer uses LibMagic.

! Empty lines in the r2agent HTTP auth file now deny all access (previously they permitted any password through).

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Wiz Code

1 RELEASE · 2026-08-13

SAST / SCA / Code Security

Wiz Code is a developer security tool that scans code and infrastructure-as-code for vulnerabilities and misconfigurations during the development lifecycle.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 4/5

Validates whether exposed secrets are actually live via a Wiz Red Agent that actively tests what a leaked credential can reach, and inventories developers' personal public GitHub repos to close the unmanaged-secret visibility gap, mapping blast radius on the Security Graph with the exact commit and pusher for owner-assigned remediation.

> Wiz Code extends secret exposure detection beyond organizational boundaries to developers' personal public GitHub repositories, pairing this with AI-driven exploitability validation and full attack-path visibility on the Security Graph.

└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

01 Identity-attached findings and remediation workflow integration

NEW

67

Surfaces identity-attached findings with the exact commit, GitHub user, and email that pushed a secret, enabling immediate owner-assigned remediation without manual investigation. These findings integrate into existing remediation workflows including Jira ticket creation, Slack developer notification, and tracked token rotation and revocation.

— Names concrete remediation destinations (Jira, Slack) giving a clear starting point.

2026-08-13-8ecb9077

THINNER COVERAGE BELOW

02 AI Red Agent exploitability validation for exposed credentials

NEW

58

Introduces exploitability validation via the Wiz Red Agent — an AI-powered attacker that actively tests whether an exposed credential is live and what it can reach, going beyond traditional validity checks.

— Names the Red Agent mechanism but no trigger or config surface given.

2026-08-13-8ecb9077

03 Security Graph attack path mapping for personal-repo secrets

NEW

54

Enriches secrets found in personal repositories onto the Wiz Security Graph to show the full attack path and blast radius of a compromised credential, including downstream resource access.

— Names the Security Graph mechanism but no navigation steps given. 2026-08-13-8ecb9077

04 **Personal GitHub repository secret inventory** NEW 45

Wiz Code adds automated inventory and correlation of developer identities to their personal public GitHub repositories outside org boundaries, closing the visibility gap where unmanaged secret exposure lives.

— Describes the capability but no concrete config or command shown. 2026-08-13-8ecb9077

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 Tenzir

2 RELEASES · 2026-08-07 → 2026-08-10

Detection Engineering

Tenzir Changelog is an RSS feed tracking updates, new features, and fixes for Tenzir's open-source security data pipeline and threat detection platform.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

Adds native beacon-detection analytics to TQL with cadence (deltas, mad, skewness) and periodicity (autocorrelation, periodogram, dominant_period) functions for identifying periodic C2 beaconing directly in pipelines, plus automatic NetFlow detection and richer ClickHouse JSON handling.

> Tenzir shipped native TQL functions for beacon and anomaly detection alongside ClickHouse and NetFlow ingestion improvements, while the Platform added pipeline ID search and made telemetry collection opt-in.

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

01 **Beacon detection analytics functions in TQL** NEW 86

Adds cadence analysis functions `deltas`, `mad`, and `skewness` and periodicity detection functions `autocorrelation`, `periodogram`, and `dominant_period` to TQL, providing research-grade tooling for beacon and anomaly detection and for identifying periodic C2 beacon patterns.

Detect beacon periodicity in connection logs by computing the dominant period of inter-arrival deltas.

```
$ from connections | summarize times=collect(ts) by src_ip,
dest_ip | extend d=deltas(times) | extend
period=dominant_period(autocorrelation(d))
```

Identify irregular beaconing candidates by flagging sources whose inter-arrival skewness or MAD exceeds expected thresholds.

```
$ from connections | summarize times=collect(ts) by src_ip |
extend d=deltas(times) | where mad(d) > 5s or skewness(d) > 2.0
```

— Named functions plus two runnable pipeline examples

Tenzir Node v6.10.0: Native time-series …

THINNER COVERAGE BELOW

02 **JSON column handling in to_clickhouse** IMPROVED 52

Adds non-top-level JSON column support and automatic JSON column detection to `to_clickhouse`, including detection after `ocsf::cast`.

— Names function and cast but gives no example Tenzir Node v6.10.0: Native time-series …

03 **Pipeline search by ID** IMPROVED 45

The pipeline list search box now matches on pipeline ID, letting you paste an ID to jump directly to a pipeline.

— Clear UI path but no exact navigation or config

Tenzir Platform v1.38.0: Opt-in telemetr…

04 **Opt-in telemetry consent** IMPROVED 40

Telemetry collection is now opt-in, gated behind explicit user consent via the Account page.

— Names UI location but no further mechanism detail

Tenzir Platform v1.38.0: Opt-in telemetr…

05 **Automatic NetFlow format detection** IMPROVED 25

Enables automatic NetFlow format detection for simplified ingestion pipelines.

— Bare description, no mechanism or example given

Tenzir Node v6.10.0: Native time-series …

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · August 14, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe