THE CYBER TOOLCHAIN NO. 006

Head

THE WEEK'S TOP PICKS

PUBLISHED AUGUST 8, 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	**Dreadnode**	AI-Powered Offensive Security
02	**Metasploit Pro**	Exploitation & C2
03	**Cribl**	Security Data Pipeline
04	**Aikido Security**	Cloud Security Posture
05	**Prowler**	Cloud Security Posture

01 Dreadnode

2 RELEASES · 2026-08-04

AI-Powered Offensive Security

Dreadnode is an AI-powered security research platform that helps teams build, run, and evaluate offensive AI agents for penetration testing and red teaming.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Adds a full traditional-ML red teaming suite — evasion, extraction, membership inference, and model inversion (20+ named attacks across tabular, image, and text) runnable via a new `dn airt run-classifier` CLI directly against a predict endpoint with no SDK code, plus MITRE ATLAS/SAIF/NIST mappings. The ATLAS multi-agent attack framework and per-finding privacy metrics round out coverage few AI red team platforms match.

> Dreadnode shipped a full traditional-ML red-teaming suite — 20+ attacks runnable via the new `dn airt run-classifier` CLI — and an ATLAS multi-agent attack framework, alongside an overhauled Agent Output triage workspace, new agent-memory tooling, air-gapped hub-sync commands, and five new data-integration connectors.

└──▷ WHAT SHIPPED · 20 FEATURES most completely described first ? what's the number?

01 **Traditional-ML attack suite in AIRT** NEW 88

Adds `dn airt run-classifier` CLI command to run classic-ML attacks directly against a predict endpoint without SDK code, covering four traditional-ML attack families in AIRT: evasion, extraction, membership inference, and model inversion. Includes 20+ named attacks (equation-solving, Jacobian, copycat, knockoff, PWWS, BAE, TextBugger, threshold, label-only, among others) across tabular, image, and text modalities, with per-finding adversarial trajectory traces and MITRE ATLAS/SAIF/NIST compliance mappings. Findings carry richer metrics: extraction reports KL divergence and accuracy retention, membership inference reports attack accuracy, precision, and recall, and model inversion reports reconstruction quality from black-box query-based reconstruction of class-representative inputs.

— Names CLI, attack families, metrics; lacks a runnable example

Traditional-ML red teaming suite — evasi…

02 **Configurable project-memory preload limit** IMPROVED 84

New `--project-memory-preload-limit` flag sets the number of memories preloaded at session start (range: 1–200), replacing a hardcoded default of 20.

Increase the number of project memories available at session start beyond the default of 20 to give agents broader context on large engagements.

```
$ dn ... --project-memory-preload-limit 100
```

— Exact flag, range, default, and usage example `Traditional-ML red teaming suite — evasi…`

03 **Air-gapped hub sync via `dn capability sync` and `dn task sync`** NEW 83

Adds `dn capability sync` and `dn task sync` CLI commands to pull hub content from SaaS (or another on-prem instance) into a local directory and push it to a self-hosted instance, enabling air-gapped on-prem loading.

— Names both commands and their direction of sync

`Agent memory, triage, and session monito…`

04 **Agent selection flag in evaluation runs** NEW 78

New `dn evaluation create --agent <name>` flag specifies which agent runs in an evaluation; eval output reports which agent actually ran, aiding reproducibility.

Pin a specific agent to an evaluation run so results are reproducible and the run report identifies which agent executed.

```
$ dn evaluation create --agent <name>
```

— Exact flag and example command given `Traditional-ML red teaming suite — evasi…`

05 **`/report-bug` command in the TUI** NEW 72

Adds `/report-bug` command (or `r` in console) to the TUI, assembling a privacy-reviewed diagnostic bundle and submitting it without losing logs after a crash.

File a bug report from the TUI immediately after a crash without losing session logs.

```
$ /report-bug
```

— Exact command and shortcut with example `Agent memory, triage, and session monito…`

06 **Agent Output findings triage workspace** IMPROVED 70

Agent Output is now a severity-first triage workspace: findings can be searched, sorted, and filtered, and each finding can be exported to CSV, deep-linked, or copied into a report. Finding status and notes can now be edited directly from a finding's detail panel. On the Agent Outputs page, the 'Needs review' and 'Critical + High' stat cards are now clickable filter toggles, and the items grid gained a sortable 'Updated' column.

— Names UI controls and export format, no API endpoint

Agent memory, triage, and session monito… | Traditional-ML red teaming suite — evasi…

07 **Session monitoring redesign in Agent Sessions** IMPROVED

68

Adds an 'Inactive' filter to Agent Sessions to isolate sessions idle past 15 minutes but not yet archived. Adds a 'Last Msg' relative-activity column and highlights sessions active in the last 15 minutes. Adds forward/backward transcript pagination with tailing mode. The sessions group-by control (None / Workflow / Capability / Agent / Source) is now available to all orgs, with workflow runs collapsed by default.

— Names filters and controls, no exact UI path | Agent memory, triage, and session monito…

08 **Web-security runtime skills: recon and static analysis** NEW

62

New web-security initial-recon skill provides structured reconnaissance for domain, wildcard, ASN, IP-list, and CIDR entry points using the ProjectDiscovery toolchain. Adds AST-based JS/TS static analysis via `ast-grep` in the web-security runtime, removing the need for semgrep, CodeQL, or a jxscout license.

— Names toolchain and tool but no exact command | Agent memory, triage, and session monito… | Traditional-ML red teaming suite — evasi…

THINNER COVERAGE BELOW

09 **AIRT demo notebooks** NEW

58

Five self-contained AIRT demo notebooks cover model evasion, extraction, membership inference, jailbreaks, and multi-agent attacks — runnable via `dn login` with no provider keys required.

— Names command to run, but not the notebooks themselves

Traditional-ML red teaming suite — evasi…

10 **ATLAS multi-agent attack framework** NEW

55

Adds the ATLAS multi-agent attack framework for running adaptive campaigns against deployed multi-agent environments across three attack surfaces, with executed tool calls captured per-trial and surfaced in findings across the UI, API, and SDK.

— Describes scope and surfaces but no concrete command

Traditional-ML red teaming suite — evasi…

11 **`llm-wiki` and `analysis-ledger` capabilities** NEW

54

Adds the `llm-wiki` capability (compiles source material into linked wiki pages) and `analysis-ledger` capability (maintains a resumable code-analysis frontier with evidence-bearing claims).

– Names both capabilities and their function, no command shown

Agent memory, triage, and session monito…

12 **Executive/CISO summary panel** NEW 52

New executive/CISO summary panel on the assessment view shows attack-surface posture, query cost to compromise, risk score, and a plain-English headline for stakeholder reporting.

– Names panel contents, no navigation path given Traditional-ML red teaming suite – evasi…

13 **Memories page under Agents** NEW 52

New Memories page under Agents lets operators view, inspect, and manage project memory entries, including version history and diffs.

– Names page location, no config detail Agent memory, triage, and session monito…

14 **Table view added to Hub browse pages** IMPROVED 52

Models and Datasets Hub browse pages add a table view alongside cards, with a unified toolbar offering separate sort and filter controls across all four Hub browse pages.

– Names view and scope, no navigation path Agent memory, triage, and session monito…

15 **Agent Output item recovery from span backups** IMPROVED 48

Agent Output items from failed `report_item` POSTs are automatically recovered from span backups when a session is frozen.

– Names failure mode and recovery mechanism, no manual trigger

Agent memory, triage, and session monito…

16 **AI Red Teaming Learning Guide** NEW 45

New AI Red Teaming Learning Guide adds eight documentation pages covering all four traditional-ML attack categories and generative AI attacks, with runnable SDK, TUI, and CLI examples.

– Names page count and coverage, no direct link Traditional-ML red teaming suite – evasi…

17 **Evaluation runtime health recording** NEW 43

Evaluation runtime now records the health of all runtime components (broken tools, skills, policies) at run time, so degraded capabilities are visible after the fact.

– Explains what's recorded, no query interface named

Traditional-ML red teaming suite – evasi…

18 **Data-integration connectors for Atlassian, Azure DevOps, GitHub, GitLab, and Linear**

NEW

40

Adds five data-integration connectors — Atlassian, Azure DevOps, GitHub, GitLab, and Linear — as platform capabilities.

— Names all five connectors, no setup steps given

Agent memory, triage, and session monito…

---

19 **Organization picker in new workspace dialog** IMPROVED

37

New workspace dialog includes an organization picker to create a workspace in any org without switching context.

— Brief description of dialog change Agent memory, triage, and session monito…

---

20 **Media gallery for session tools** NEW

33

New dedicated media gallery collects all images and other media logged by a session's tools in one place for reviewers.

— Brief description, no location or size given Traditional-ML red teaming suite — evasi…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Metasploit Pro

1 RELEASE · 2026-08-03

Exploitation & C2

Metasploit Pro is a commercial penetration testing framework for identifying, exploiting, and validating security vulnerabilities in systems.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

Adds a `MALLEABLEC2` option to every Meterpreter HTTP(S) payload to reshape C2 traffic with standard Malleable C2 profiles, alongside an HTTP-to-SMB NTLM relay module, eight new MCP tools for module/session control, and RISC-V payloads and encoders. A broad offensive release that meaningfully expands evasion and automation for red teamers.

> Metasploit Pro 5.1.0-2026073001 adds Malleable C2 traffic shaping for Meterpreter payloads, eight new MCP server tools wrapping RPC endpoints for module and session control, expanded CertificateTrace and KerberosTicketTrace visibility, RISC-V and Windows-on-ARM payload support, and a new HTTP-to-SMB NTLM relay module.

└──▷ WHAT SHIPPED · 13 FEATURES most completely described first ? what's the number?

01 MCP server tools and RPC integration NEW

95

Adds eight new MCP tools wrapping existing RPC endpoints for module and session control, backed by strict per-tool input validation, plus a `--enable-dangerous-actions` flag to the `msfmcpd` CLI to opt in to dangerous actions across those tools. Adds `run_uuid` job tracking across all module types when executed through RPC endpoints to support MCP tool integrations, and Bearer Authentication support to the MCP server plugin. Also adds `resource` and `parents` fields to the `msf_service_info` MCP tool, `resource` to `msf_vulnerability_info`, `data` to `msf_note_info`, and `realm_key` / `realm_value` to `msf_credential_info`.

Start the MCP server with dangerous module-execution actions enabled so MCP tool integrations can run and check modules via RPC.

```
$ msfmcpd --enable-dangerous-actions
```

— Names the CLI flag, tool count, and every added field August 3, 2026: 5.1.0-2026073001

02 Malleable C2 traffic shaping for Meterpreter payloads NEW

90

Adds a `MALLEABLEC2` option to all Meterpreter HTTP(S) payloads (Windows, Java, Python, PHP, Linux) to load a standard Malleable C2 profile and reshape HTTP(S) traffic.

The Pro edition adds support for selecting Malleable C2 profiles from the single module run page, the payload generator, and the global persistent listeners page.

Shape Meterpreter HTTPS traffic using a Malleable C2 profile to blend into expected network baselines during an engagement.

```
$ use exploit/multi/handler
set PAYLOAD windows/x64/meterpreter/reverse_https
set LHOST 10.0.0.1
set LPORT 443
set MALLEABLEC2 /opt/profiles/jquery.profile
run
```

— Named option, payload types, Pro pages, and a runnable example

August 3, 2026: 5.1.0-2026073001

03 **CertificateTrace expanded across protocols** IMPROVED 78

Extends `CertificateTrace` peer cert tracing (using the existing `off` / `metadata` / `full` enum) to LDAP over TLS, RDP, and PostgreSQL over TLS (passing the `SSL` datastore option through to the underlying Postgres connection), and to HTTP modules connecting over HTTPS. Adds a Certificate Signing Request (CSR) trace so operators can see the CSR printed when requesting certificates from AD CS.

— Names enum values and every extended protocol, no example

August 3, 2026: 5.1.0-2026073001

04 **KerberosTicketTrace output granularity** IMPROVED 60

Adds improved granularity to `KerberosTicketTrace`: operators can now print the full trace, only tickets, or only metadata. Reuses `KerberosTicketTrace` formatting for offline Kerberos artifact workflows so stored and converted tickets render with the same output style as live authentication.

— Describes modes but no exact flag syntax given August 3, 2026: 5.1.0-2026073001

05 **FTP fingerprinting methods in Ftp mixin** NEW 60

Adds `ftp_fingerprint` and `ftp_list_directory` methods to the `Exploit::Remote::Ftp` mixin, leveraging recog to fingerprint targets from their banners.

— Names mixin and methods but no usage example August 3, 2026: 5.1.0-2026073001

06 **RISC-V payload and encoder support** NEW 60

Adds four encoder variants for both RISC-V 32-bit and 64-bit little-endian architectures, plus `reverse_tcp` and `bind_tcp` stagers and a shell command stage for RISC-V 64-bit and 32-bit little-endian Linux targets.

— Names architectures and stager types, no example | August 3, 2026: 5.1.0-2026073001

07 **Apache htaccess web shell persistence module** NEW 60

Adds a new persistence module `exploits/linux/persistence/apache_htaccess` that plants a mod_cgi `.htaccess` web shell on Linux Apache targets.

— Names exact module path | August 3, 2026: 5.1.0-2026073001

THINNER COVERAGE BELOW

08 **HTTP-to-SMB NTLM relay module** NEW 55

Adds an HTTP-to-SMB NTLM relay server module that relays incoming NTLM HTTP authentication requests to multiple SMB servers to establish SMB sessions.

— Describes mechanism without exact module path | August 3, 2026: 5.1.0-2026073001

09 **FTP version scanner reports bannerless hosts** IMPROVED 55

Improves `scanner/ftp/ftp_version` to report the service when a host is up but no FTP banner is received.

— Names exact module and the behavior change | August 3, 2026: 5.1.0-2026073001

10 **Payload generator warning for risky encoders** IMPROVED 50

Pro edition adds a warning alert in the payload generator when `encoder/none` or `encoder/eicar` is selected.

— Names exact encoders that trigger the warning | August 3, 2026: 5.1.0-2026073001

11 **Linux fetch payloads and fetch/TFTP servers** NEW 45

Adds Linux fetch multi payloads, a fetch server for FTP-based fetch payloads, and a TFTP server to `rex/proto`.

— Names rex/proto addition but no usage shown | August 3, 2026: 5.1.0-2026073001

12 **Windows on ARM reverse-TCP shell payload** NEW 30

Adds a Windows on ARM (AArch64) reverse-TCP command-shell payload.

— Bare feature name with no further mechanism | August 3, 2026: 5.1.0-2026073001

13 **Network topology page enhancements** IMPROVED 15

Pro edition adds multiple enhancements to the network topology page.

— No specifics given about the enhancements August 3, 2026: 5.1.0-2026073001

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Cribl

2 RELEASES · 2026-07-31

Security Data Pipeline

Cribl is an observability pipeline platform that helps organizations collect, process, route, and transform machine data from any source to any destination.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

Launches an AI Observability app giving one console to search and investigate AI telemetry, with Sensitive Data Traces that pinpoint exactly where PII, credentials, and API keys appear in any trace or session, plus tool-call audit and cost/usage dashboards. A complete capability for teams trying to see and govern what their AI is actually doing.

> Cribl launched a new App for AI Observability, giving teams a single place to search, monitor cost/token usage, detect sensitive data exposure, and investigate individual AI traces and sessions across models and providers.

└──▷ WHAT SHIPPED · 9 FEATURES most completely described first ? what's the number?

01 Sensitive Data Traces for PII and credential exposure NEW

80

Sensitive Data Traces pinpoint exactly where PII, credentials, and API keys appear within a trace or session, with analysis by category, user, model, and tool over time.

Surface PII or credential leakage in AI sessions by reviewing Sensitive Data Traces — useful during a security review or incident triage.

> 📍 In Cribl.Cloud, open the AI Observability app › Security and Compliance Dashboard › Sensitive Data Traces, then filter by subcategory (PII, credentials, API keys) and by user or model to locate the exact trace or session where sensitive data appeared.

– Exact dashboard path and filter steps to locate exposures

What is your AI actually doing? Introduc… What is your AI actually doing? Introduc…

02 Sessions and Traces investigation view NEW

80

The Sessions and Traces view provides end-to-end trace visibility with filterable session lists, span timelines, and full prompt, response, and tool-call payloads.

Investigate a model quality regression by jumping from a dashboard anomaly directly into the full prompt and response payloads for that session.

> 📍 In Cribl.Cloud, open the AI Observability app › Sessions and Traces, filter the session list by model and time range, select the suspect session, and inspect the

span timeline alongside full prompt, response, and tool-call payloads.

— Exact navigation and filtering steps to inspect a session

What is your AI actually doing? Introduc… | What is your AI actually doing? Introduc…

03 **Goals and Monitors for AI budget tracking** NEW 70

Goals and Monitors continuously evaluate live AI telemetry against user-defined targets, such as month-to-date token spend versus budget, alerting when thresholds are crossed.

Track AI spend against budget targets to catch token overruns before they become a finance problem.

📍 In Cribl.Cloud, open the AI Observability app › Goals and Monitors, create a new monitor targeting month-to-date token spend, and set the budget threshold for your team or model.

— Exact UI path and steps to create a budget monitor given

What is your AI actually doing? Introduc… | What is your AI actually doing? Introduc…

THINNER COVERAGE BELOW

04 **Usage and Cost Dashboard for AI spend** NEW 50

The Usage and Cost Dashboard surfaces spending trends, token consumption, top users and models, LLM efficiency anti-patterns, and an activity heatmap. Accompanying Usage Efficiency tooling exposes LLM efficiency issues to help improve token efficiency and model effectiveness.

— Lists specific metrics shown but no navigation steps given

What is your AI actually doing? Introduc… | What is your AI actually doing? Introduc…

05 **Cribl App for AI Observability launched** NEW 45

A new Cribl app provides a single platform to search, investigate, and report on AI telemetry across models, tools, and environments. It is built on Cribl's AI Platform for Telemetry Lakehouse engine, enabling data collection, transformation, storage, and flexible analysis via search or dashboards.

— Names the underlying engine but gives no concrete entry point

What is your AI actually doing? Introduc… | What is your AI actually doing? Introduc…

06 **Security and Compliance Dashboard** NEW 45

The Security and Compliance Dashboard tracks sensitive data exposure rates, findings by type and user, tool-call audit trails, and content risk indicators.

— Names tracked metrics but no direct navigation example

What is your AI actually doing? Introduc… What is your AI actually doing? Introduc…

07 **AI Insights conversational agent** NEW 40

AI Insights offers a conversational AI agent for querying usage, cost, and security data across the organization in plain language.

— Describes capability but no example query or path

What is your AI actually doing? Introduc… What is your AI actually doing? Introduc…

08 **AI-generated anomaly briefs** NEW 25

AI-generated briefs automatically surface anomalies and areas needing attention across usage, cost, and security data, without requiring manual investigation.

— Vague on trigger conditions or where briefs appear

What is your AI actually doing? Introduc… What is your AI actually doing? Introduc…

09 **Org View spend roll-up** NEW 20

Org View delivers a hierarchical roll-up of spend, tokens, and requests across the organization.

— One-line description with no mechanism or usage path

What is your AI actually doing? Introduc… What is your AI actually doing? Introduc…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Aikido Security

1 RELEASE · 2026-08-04

Cloud Security Posture

Aikido Security is a developer-focused security platform that scans code, dependencies, containers, and cloud infrastructure to detect and prioritize vulnerabilities for fast remediation.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

Adds code-based DSPM that uses LLM agents to trace sensitive data from sources to sinks entirely from code — no production data store connection required — flagging credential leakage in URLs, API over-exposure, unsafe logging, and over-broad third-party/AI sharing. A novel path to data-exposure mapping without touching live data.

> Aikido introduces code-based DSPM with a new Data Exposure Audit that traces sensitive data flows through source code to find exposure risks without ever connecting to production data stores.

└──▷ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

01 **Data Exposure Audit for code-based DSPM** NEW 75

A new Data Exposure Audit feature uses LLM agents to trace sensitive data from sources (request bodies, uploads, DB reads, config, application logic) to sinks (logs, third-party APIs, LLM prompts, caches, queues, storage) entirely from code, with no production data store connection required. It detects credential and token leakage in URLs, query strings, redirects, and clone URLs; identifies API over-exposure where endpoints return more fields than the caller is entitled to or lack tenant ownership checks; flags unsafe logging and exports of unredacted PII and secrets in logs, traces, caches, and dumps; surfaces over-broad third-party and AI data sharing to vendors, analytics, LLM prompts, and vector stores; detects weak storage patterns such as regulated data stored unencrypted or protected only with base64, MD5, or SHA-1; and finds incomplete deletion flows that leave downstream records, caches, shadow copies in staging, and snapshots intact. Each Data Exposure Audit consumes credits from the Aikido Wallet.

— Rich mechanism and scope, but no named flag, endpoint, or config key.

New - DPSM: Find data exposure risks wit…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 Prowler

2 RELEASES · 2026-08-03 → 2026-08-06

Cloud Security Posture

Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 3/5

Adds 20 AWS privilege-escalation Attack Path queries from pathfinding.cloud (with 22 new escalation combinations in the privesc check), plus 12 new M365 CIS checks, a new Huawei Cloud provider, and expanded MCP tooling for scans, triage, and alerts. A substantial multi-cloud posture and attack-path release.

> Prowler added a new Huawei Cloud provider and 11 AWS Nitro Enclaves checks, expanded Attack Paths with dozens of AWS privilege-escalation queries and 12 M365 compliance checks, and extended its MCP Server with new integration, user, role, alerting, and scheduling tools alongside SAML multi-domain SSO, compliance watchlists, and a consolidated cross-provider compliance view.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 Huawei Cloud provider support NEW

90

Adds a new `huaweicloud` provider authenticated via `HUAWEICLOUD_ACCESS_KEY_ID` and `HUAWEICLOUD_SECRET_ACCESS_KEY` environment variables, run with `prowler huaweicloud`. It ships 25 checks covering CTS, ECS, ELB, EVS, IAM, KMS, OBS, RDS, VPC, and WAF, plus the CIS Huawei Cloud Foundations Benchmark 1.0.

Scan a Huawei Cloud account for the first time using IAM user access key credentials.

```
$ export HUAWEICLOUD_ACCESS_KEY_ID="your-access-key-id"
export HUAWEICLOUD_SECRET_ACCESS_KEY="your-secret-access-key"
prowler huaweicloud
```

— Exact env vars, command, and example provided 5.37.0

02 MCP Server integrations, users, and roles tools NEW

80

Extends the Prowler MCP Server (Cloud and self-hosted Local) with three tool families: Integrations (full lifecycle for Amazon S3, AWS Security Hub, and Jira — create, update, re-check, delete, and dispatch findings as Jira work items), Users (list tenant users and identify the authenticated user), and Roles (browse RBAC roles, inspect capabilities, and set a user's role). A new `prowler_cloud_*` namespace of 32 Cloud-only MCP tools, exposed at `https://mcp.prowler.com/mcp`, covers Alerts (create/manage alert rules,

dry-run conditions, browse fired-alert history), Findings Triage (set triage status and attach notes), Scan Scheduling (daily, interval, weekly, or monthly recurring scans per provider or bulk), and Scan Configurations (build reusable configs and attach to providers).

— Names tool families, namespace, and endpoint URL `5.37.0`

---

03 **Attack Paths AWS privilege-escalation query library** NEW

75

Attack Paths gained 20 new AWS privilege-escalation detection queries from pathfinding.cloud covering `iam:PassRole` escalations across AWS Batch, Braket, Cognito Identity, ECS, EMR, EMR Serverless, GameLift, Glue, EC2 Image Builder, Kinesis Analytics, HealthOmics, EventBridge Scheduler, SSM, and Step Functions, plus CodeDeploy/Step Functions existing-resource abuse, permissions-boundary removal with role assumption, and IAM Identity Center permission-set policy injection. It also added four named queries — `STS-002` (cross-account role trust), `STS-003` (wildcard role trust), `IAM-022` (user permissions-boundary removal), and `SSO-001` (IAM Identity Center permission-set escalation).

— Names specific queries and services but no run command `5.38.0` `5.37.0`

---

04 **iam_policy_allows_privilege_escalation check expanded** IMPROVED

70

The `iam_policy_allows_privilege_escalation` check gained 22 additional escalation combinations sourced from pathfinding.cloud, covering AWS Batch, Braket, CodeDeploy, Cognito Identity, ECS, EMR, EMR Serverless, GameLift, Glue, EC2 Image Builder, Kinesis Analytics, HealthOmics, and EventBridge Scheduler.

— Names check and services, no example command given `5.38.0`

---

05 **AWS Nitro Enclaves CSPM checks** NEW

70

Adds 11 checks for AWS Nitro Enclaves: five `ec2_confidential_workload_host_*` checks for parent-instance hygiene (IMDSv2, public IP, ingress, vsock proxy ports, host state) and six `kms_key_enclave_*` checks for KMS attestation-policy correctness, mapped across 23 compliance frameworks including NIST 800-53 Rev 5, PCI-DSS v4.0, ISO 27001:2022, SOC 2, HIPAA, and MITRE ATT&CK.

— Names check patterns and frameworks, no run example `5.37.0`

---

06 **Consolidated compliance view across providers** IMPROVED

65

The Compliance view in Prowler Cloud now groups frameworks across all providers of the same type, rolling up into a consolidated posture with per-provider breakdown, findings drill-down, and a combined executive PDF report. Tabs are renamed to `Single Scan` and `Multiple Scans`, with `Multiple Scans` as the new default.

— Explains mechanism and tab names, no deep steps 5.37.0

07 **12 new M365 compliance checks** NEW 60

Added `admincenter_shared_bookings_disabled`, `defender_priority_account_protection_enabled`, `defender_strict_preset_security_policy_enabled`, `exchange_owa_mailbox_policy_personal_accounts_disabled`, `exchange_organization_reject_direct_send_enabled`, `teams_external_access_trial_tenants_blocked`, `entra_device_registration_join_restricted`, `entra_device_registration_max_devices_per_user_limited`, `entra_device_registration_global_admins_not_local_admins`, `entra_device_registration_registering_user_not_local_admin`, `entra_device_registration_laps_enabled`, and `entra_policy_default_user_cannot_read_bitlocker_keys` checks for the M365 provider, mapped to CIS Microsoft 365 Foundations Benchmark v7.0.0 controls 1.3.9, 2.4.1, 2.4.2, 6.3.2, 6.5.5, 8.2.4, and the 5.1.4.x device-registration family.

— Names all 12 checks and controls, no run example 5.38.0

08 **Secrets and encryption checks across AWS services** NEW 60

Adds `codecommit_repository_no_secrets` (with a new `codecommit` service) to scan files at the tip of each repository's default branch for hardcoded secrets, `glue_catalog_connection_no_secrets` to detect secrets in Glue Data Catalog connection properties, and `sagemaker_endpoint_config_kms_encryption_enabled` to verify SageMaker endpoint configurations use a KMS key for storage volume encryption.

— Names three checks, no example command shown 5.37.0

09 **OCSF findings enriched with MITRE ATT&CK** IMPROVED 60

Populates `finding_info.analytic` and `finding_info.attacks` in OCSF detection finding output with Prowler check rules and MITRE ATT&CK technique/tactic objects, for findings that carry MITRE ATT&CK compliance metadata.

— Names OCSF fields, condition described, no example 5.37.0

THINNER COVERAGE BELOW

10 **Attack Paths query metadata and Hub links** IMPROVED 57

Attack Paths query metadata on the queries endpoint now exposes an outcome category — Code execution, Privilege escalation, Public exposure, or Resource inventory — enabling terminal outcome nodes in the graph. The query info panel also now links every query to its page on Prowler Hub.

— Explains mechanism but no exact UI/API path given 5.37.0 5.38.0

11 **Compliance Watchlist for pinned frameworks** NEW 55

Compliance Watchlist (Prowler Cloud/Private Cloud) lets teams pin frameworks from any compliance view, filter the Compliance section to show only pinned frameworks, and see pinned framework scores on the Overview page.

— Explains mechanism, UI path only, no exact steps 5.38.0

12 **Container build supply-chain hardening** IMPROVED 55

Prowler API, UI, SDK, and MCP container images now publish per-architecture SBOMs and build-provenance attestations as OCI attestations. Grype now runs alongside Trivy across container-image security gates, blocking on fixable high and critical findings, and API/SDK builds now verify checksums of downloaded PowerShell, Trivy, and zizmor binaries before installation.

— Names tools and process, not user-actionable 5.38.0

13 **Guided GCP Organization onboarding** IMPROVED 55

GCP Organization onboarding in Prowler Cloud is now a single guided flow: provide an org-level credential, select folders and projects from a hierarchy tree, set aliases, test the connection, and register each project as a provider.

— Describes flow steps, no exact UI navigation given 5.37.0

14 **SAML multi-domain SSO** IMPROVED 50

SAML SSO now supports one primary email domain plus up to 19 additional domains under a single Identity Provider configuration, using a single stable ACS URL (Prowler Cloud/Private Cloud).

— Gives numbers but no configuration steps shown 5.38.0

15 **Context-aware Lighthouse AI** IMPROVED 50

Lighthouse AI (Prowler Cloud/Private Cloud) messages now carry the current page, open finding or resource, and its metadata, and Lighthouse can invoke all MCP tool families — scan configurations, scheduling, finding triage, alert rules, integrations, users, and roles — subject to RBAC.

— Explains mechanism, no exact commands given 5.37.0

16 **Sign-in method tags in Users table** NEW 40

The Users table in Prowler Cloud now shows each account's sign-in methods as tags, including email/password, Google, GitHub, SAML with linked domains, and Partner SSO.

— Describes UI change only, no mechanism depth 5.38.0

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · August 8, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe