THE CYBER TOOLCHAIN

NO. 018

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 17, 2026 · EVERY WEEKDAY

The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 28 tools

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

ProjectDiscovery dnsx ⓘ v1.3.0 NOTES ↗

dnsx is a fast and multi-purpose DNS toolkit allow to run multiple DNS queries of your choice with a list of user-supplied resolvers.

dnsx v1.3.0 adds automatic wildcard detection, custom output templates, query timing, and inline host input.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.3.0
https://github.com/projectdiscovery/dnsx.git
# already have the repo? check out this version:
$ git checkout v1.3.0
```

└──▷ TRY IT

Automatically filter wildcard domains during subdomain brute-force without manual configuration.

```
$ dnsx -d example.com -w wordlist.txt -auto-wildcard -silent
```

Emit a custom one-liner per result (host, A record, and query time) to feed into downstream parsers or reports.

```
$ subfinder -silent -d example.com | dnsx -silent -a -ot '{{host}}
{{a}} {{query-time}}'
```

Resolve a quick ad-hoc list of hosts passed inline without creating a file.

```
$ dnsx -l example.com,api.example.com,dev.example.com -a -resp
```

- Adds `-auto-wildcard` flag for automatic wildcard DNS detection, supporting multiple registrable roots and A/AAAA-based filtering.
- Adds `-output-template` / `-ot` flag for customizable output using template fields such as `{{host}}`, `{{a}}`, `{{aaaa}}`, `{{ip}}`, and `{{query-time}}`.
- Adds `query-time` response field to expose per-query DNS timing in output.

› Supports inline and comma-separated host input directly on the command line (no file required).

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Promptfoo ⓘ snapshot-20260717 NOTES ↗

Promptfoo is an open-source tool for testing, evaluating, and red-teaming LLM applications to detect vulnerabilities, compare outputs, and ensure AI reliability.

Promptfoo snapshot-20260717 adds adaptive rate limiting, Transformers.js local inference, telecom red team plugins, and video generation providers.

└──▷ TRY IT

Load multiple environment files (e.g., base secrets plus environment-specific overrides) in a single eval run.

```
$ promptfoo eval --env-file .env.base --env-file .env.staging --
prompts prompts.yaml --tests tests.yaml
```

Cap test case generation per strategy to keep red team scans fast during development.

```
$ promptfoo redteam generate -d 'telecom-smoke-test' --config
redteam.yaml
```

- Adds Transformers.js provider to run Hugging Face models locally in Node.js or the browser.
- Adds Vercel AI Gateway and Cloudflare AI Gateway providers for request routing.
- Adds AWS Bedrock Video (Nova Reel, Luma Ray 2) and Azure AI Foundry Video (Sora) for video generation.
- Adds native session endpoint support on the HTTP provider for stateful conversations.
- Adds xAI Voice provider overrides for `apiBaseUrl`, `websocketUrl`, and function call support.

+22 more

WAS THIS USEFUL? 👍 👎

Portkey ⓘ snapshot-20260717 NOTES ↗

Portkey is an AI gateway and LLMOps platform that helps developers manage, monitor, and optimize calls to large language model APIs with reliability and observability.

Portkey adds vault-backed credentials, weekly rate limits, cross-cloud GCS logging, and four new guardrail integrations in March 2026.

└──▷ HOW TO FIND IT

Pull LLM provider credentials from your existing vault at runtime instead of storing keys in Portkey directly.

> 📍 In the Portkey console, go to Settings › Integrations (or Virtual Keys) and add a secret_mappings entry pointing to your AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault secret ARN/path. Portkey fetches the value at request time.

Route Zscaler AI Guard policy checks across all LLM traffic to enforce your org's existing detection policies on inputs and outputs.

```json
{
  "guardrails": [
    {
      "type": "zscaler",
      "beforeRequestHook": true,
      "afterRequestHook": true,
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    }
  ]
}
```

Ship gateway logs from AWS to GCS using Workload Identity Federation — no long-lived GCP keys required.

```
$ GCP_WIF_AUDIENCE=<your-wif-audience>
GCP_WIF_SERVICE_ACCOUNT_EMAIL=<your-service-
account>@<project>.iam.gserviceaccount.com
# Set log storage mode to gcs_assume in your gateway config
```

› New Secret References let you point Portkey at credentials stored in AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault via `secret_mappings`, so keys are fetched at runtime from vaults you already control.

- Adds weekly rate and budget windows (`rpw`) so spend caps align with weekly planning cycles rather than only per-minute or monthly aggregates.
- Adds endpoint-scoped rate limits so chat completions, embeddings, and admin routes can each carry independent limits instead of a single global rule.
- GCP Workload Identity Federation log delivery: configure `GCP_WIF_AUDIENCE` and `GCP_WIF_SERVICE_ACCOUNT_EMAIL` so an AWS-hosted gateway writes logs to GCS without long-lived GCP keys (`gcs_assume` style).
- Organization admins can now include archived workspaces in analytics graphs and filters; saved filters also accept workspace slugs alongside IDs.

+9 more

WAS THIS USEFUL? 👍 👎

Veracode ⓘ Platform updates (2026-07-15) BLOG ↗

Veracode is a cloud-based application security platform that scans code for vulnerabilities, helping developers identify and fix security flaws across the software development lifecycle.

Veracode adds OAuth for XML APIs, six new built-in policies, OWASP/CWE 2025 standards, and a new Firewall Analytics dashboard.

- Supports OAuth Client Credentials authentication for Veracode XML APIs, not just REST APIs.
- Adds six new built-in security policies to the Veracode Platform, ahead of planned deprecation of recommended and transitional policies.
- Auto-Update OWASP and Auto-Update CWE Top 25 policy rules now enforce OWASP Top 10:2025 and MITRE CWE Top 25 for 2025 standards.
- Adds `collections_list` field to Finding Reporting API responses, returning collection names associated with an application.
- Adds Collections fields to the Applications explore in Veracode Analytics, with new Policy Compliance Overview dashboard visualizations.

+18 more

└──▷ BREAKING ON UPGRADE

- ! The `/healthcheck/status` endpoint has been removed and now returns a 404 error; use the `v2/principal` endpoint in the Identity REST API to check authenticated connectivity instead.
- ! Success events have been removed from the Firewall dataset in Firewall Analytics.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Bunkerity BunkerWeb ⓘ v1.6.13 NOTES ↗

Open-source and cloud-native Web Application Firewall (WAF)

BunkerWeb v1.6.13 adds global rate limiting, per-URI country check bypass, and post-challenge redirect control.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.6.13
https://github.com/bunkerity/bunkerweb.git
# already have the repo? check out this version:
$ git checkout v1.6.13
```

└──▷ USE IT

Protect a backend from being overwhelmed by capping total inbound requests across all clients before per-IP limits are evaluated.

```ini
USE_LIMIT_REQ_GLOBAL=yes
LIMIT_REQ_GLOBAL_RATE=500r/s
```

Exempt health-check and metrics endpoints from country-based blocking so monitoring systems in blocked regions still reach them.

```ini
COUNTRY_IGNORE_URI=/health /metrics\?.*
```

Send users to a welcome page after passing an antibot challenge instead of replaying their original (potentially POST) request.

```ini
ANTIBOT_SUCCESS_URI=/welcome
```

- Adds `USE_LIMIT_REQ_GLOBAL` and `LIMIT_REQ_GLOBAL_RATE` settings to cap the total aggregate request rate for a service across all clients and URLs, evaluated before per-IP and per-URL limits.
- Adds `COUNTRY_IGNORE_URI` to skip country whitelist and blacklist checks for URIs matching a list of PCRE expressions, matched against both the request path and the full

URI including query string.

› Adds `ANTIBOT_SUCCESS_URI` to redirect clients to a fixed URI after successfully completing a challenge, rather than returning them to the originally requested page.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

OpenVanta GrantGuard ⓘ v0.1.0 NOTES ↗

Audit, review, and clean up your Claude Code standing permissions.

GrantGuard v0.1.0 launches: audit and clean up risky Claude Code standing permissions from the CLI.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.1.0
https://github.com/OpenVanta/GrantGuard.git
# already have the repo? check out this version:
$ git checkout v0.1.0
```

└──▷ TRY IT

Quickly audit your personal Claude Code permissions for high-risk grants — inline API keys, credential-store reads, destructive wildcards — without modifying anything.

```
$ uv run grantguard.py audit
```

Open the interactive browser UI to review and selectively remove flagged permissions across a specific repo's Claude settings.

```
$ uv run grantguard.py ui /path/to/repo
```

- New `audit` command scans Claude Code settings files and classifies every `allow` rule by risk — inline credentials, credential-store reads, destructive wildcards, unprompted remote pushes, and overbroad wildcards.
- New `ui` command opens a local browser UI for interactive review and one-click removal of flagged permission grants.
- Supports `--scan` for shallow discovery of `.claude/settings*.json` under target directories, and `--deep-scan` for broad recursive discovery.
- New `--fix` flag on `audit` applies removals non-interactively, rewriting only the flagged rules while leaving all other settings intact.
- New `--tolerance` flag (`default` or `permissive`) controls whether overbroad wildcard rules are flagged or treated as safe.

+3 more

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sublime Security ⓘ

Platform v1.86 (Advanced Graymail Protection - Beta) (2026-07-16)

2 RELEASES · 2026-07-16 | BLOG ↗ | LATEST | STABLE

Sublime Security is an email security platform that detects and blocks phishing, malware, and other email threats using customizable, logic-based detection rules.

Sublime Security adds Advanced Graymail Protection (Beta) with ML-based detection, user preference learning, and automated remediation for M365 and Google Workspace.

.

└──▷ HOW TO FIND IT

Enable graymail user preference learning so Sublime personalizes bulk-mail handling per user based on their folder moves.

> 📍 In the console, go to Admin → Account and enable User Preferences.

Activate automated graymail remediation to automatically route and manage detected graymail without manual triage.

> 📍 In the console, activate the Rule 'Fuzzy Attack Score: Advanced Graymail Detection' and the Automation 'Remediate graymail flagged messages'.

- New dedicated Graymail view separates high-volume non-malicious mail (marketing, newsletters, subscriptions) from other detections across Microsoft 365 and Google Workspace.
- Adds bi-directional User Preferences: Sublime learns per-user sender preferences when users move messages into or out of Promotions or Junk/Spam and applies them to future messages.
- Introduces Fuzzy Attack Score, a new ML model that identifies unwanted bulk mail without relying on rigid signatures or sender-specific rules.
- New Graymail Automation enables automatic routing, triaging, and management of detected graymail via the included 'Remediate graymail flagged messages' automation.
- New built-in detection rule 'Fuzzy Attack Score: Advanced Graymail Detection' is included and auto-activated for orgs with 'Automatically Activate New Rules' enabled.

+1 more

└──▷ BREAKING ON UPGRADE

! The 'Flagged' view has been replaced; 'Needs Remediation' is its closest equivalent for reviewing messages awaiting action.

▾ **1 more release in this issue** · 2026-07-16

Platform v1.85.57 (Multi-Tenant Provisioning APIs) (2026-07-16) BLOG ↗

STABLE

Sublime Security adds four REST API endpoints to create, list, and manage child organizations and their licenses at scale.

- New `GET /v0/organizations/mine/child-organizations` endpoint lists all child organizations with license and usage data, supporting status filtering, name search, pagination, and depth traversal.
- New `POST /v0/organizations/mine/child-organizations` endpoint creates child organizations and optionally assigns licenses in a single API call.
- New `PUT /v0/organizations/mine/child-organizations/{id}/licensing` endpoint updates the license on an individual child organization.
- New bulk licensing endpoint applies license transitions to up to 200 child organizations at once, enabling automated renewals, upgrades, and offboarding.

WAS THIS USEFUL? 👍 👎

Vectra AI

v9.13 - Observe | 2 RELEASES | BLOG ↗ | LATEST | STABLE

Vectra AI is a cybersecurity platform that uses AI-driven threat detection and response to identify and prioritize attacker behaviors across hybrid cloud environments.

Vectra AI v9.13 adds expanded detections for Windows post-exploitation, HTTP attack tools, covert tunneling, credential theft, and cloud exfiltration.

- New detections for SMB-based lateral movement, malicious named pipe activity, and Cobalt Strike SMB beaconing across Windows environments.
- Expanded HTTP-based detection coverage for malicious downloads, uploads, certificate abuse, and offensive framework activity using trusted Windows utilities.
- New detection coverage for tunneling frameworks and covert channels, including web-based tunnels, public tunneling services, and custom BackConnect protocols.
- New detections for Active Directory Certificate Services abuse, suspicious AD database file transfers, and cloud storage exfiltration via common sync tools.

▾ 1 more release in this issue

> v9.13 - Platform | BLOG ↗ | STABLE
>
> Vectra AI v9.13 expands ERSPAN support with 8- and 16-core virtual Sensor options and improved encapsulated traffic performance.
>
> - Supports 8- and 16-core virtual Sensors for ERSPAN-based monitoring deployments.
> - Improved configuration handling for ERSPAN capture interface addressing simplifies virtual and physical Sensor setup.
> - Performance optimizations increase processing efficiency for encapsulated traffic workloads in ERSPAN environments.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

NVIDIA OpenShell ⓘ v0.0.86 NOTES ↗

OpenShell is the safe, private runtime for autonomous AI agents.

OpenShell v0.0.86 adds AWS STS AssumeRole refresh, network egress middleware, and TUI panel navigation.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.0.86 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.86
```

- Adds AWS STS AssumeRole refresh strategy and aws-s3 profile support for cloud provider integrations.
- Introduces supervisor middleware framework (RFC-0009) enabling extensible sandbox lifecycle hooks.
- Adds network egress middleware via the supervisor middleware layer for controlling outbound sandbox traffic.
- Enables Up/Down arrow key navigation across panels with overflow at list boundaries in the TUI.

└──▷ BREAKING ON UPGRADE

! The `openshell policy prove` command and Z3 dependency have been removed; any workflow or script invoking `openshell policy prove` will break.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler ⓘ 5.35.0 NOTES ↗ CODE ↗

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

Prowler 5.35.0 adds Lighthouse AI side chat and actions, one-step AWS Org onboarding, scan exclusions, and a redesigned sidebar.

└──▷ GET THIS VERSION

```
$ git clone --branch 5.35.0 https://github.com/prowler-cloud/prowler.git
# already have the repo? check out this version:
$ git checkout 5.35.0
```

- Adds Lighthouse AI side panel: open a context-aware AI chat from anywhere in the app without leaving the current page, with tabs to switch between finding details and the chat (Prowler Cloud/Private Cloud).
- Enables Lighthouse AI to take actions: connect/remove providers, trigger and schedule scans, update scan settings, and manage mute rules via chat, gated by RBAC (Prowler Cloud/Private Cloud).
- Adds one-step AWS Organizations onboarding: a single CloudFormation quick-create link deploys the management account role and a service-managed StackSet to every member account (Prowler Cloud/Private Cloud).
- Adds `excluded_checks` and `excluded_services` fields to scan configurations, allowing entire checks or services to be skipped per provider for faster, lower-noise scans.
- Introduces a redesigned sidebar with grouped sections for security, settings, and help, collapsible configuration entries, route-aware active states, and a responsive mobile overlay.

+2 more

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Anchore Grype ⓘ v0.116.0 NOTES ↗ CODE ↗

A vulnerability scanner for container images and filesystems

Grype v0.116.0 adds Go symbol-based reachability analysis, Ubuntu ESM support, architecture-aware matching, and Chainguard OSV integration.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.116.0 https://github.com/anchore/grype.git
# already have the repo? check out this version:
$ git checkout v0.116.0
```

- Adds lightweight Go symbol reachability analysis to suppress false positives for stdlib, golang.org/x, and third-party Go vulnerability advisories — only flags binaries that actually import the vulnerable symbols.
- Deduplicates Go matches that are aliases of each other (same CVE reported under both a govulndb GO-* ID and its GHSA).
- Supports Ubuntu ESM (Extended Security Maintenance) vulnerability matching.
- Adds a Chainguard OSV transformer for processing Chainguard vulnerability data.
- Duplicates RHSA advisories across all applicable RHEL minor versions for more complete RHEL coverage.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress ⓘ IP-Based Expected Rules Now Available in Managed ITDR

NOTES ↗

Huntress is a managed cybersecurity platform that detects, investigates, and remediates threats targeting small and mid-sized businesses through persistent footholds and endpoint monitoring.

Managed ITDR now supports IP- and CIDR-based Expected Unwanted Access rules to suppress known-good login escalations.

└──▷ HOW TO FIND IT

Suppress login escalations from a known office or remote-worker IP range without creating overly broad country or VPN exceptions.

> 📍 In the Huntress console, go to Managed ITDR › Expected Unwanted Access rules and create a new rule scoped to a single IP address or CIDR range.

› Adds IP-based Expected Unwanted Access rules in Managed ITDR, letting partners and customers suppress escalations for individual IPs or entire CIDR ranges without disabling other detections.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Tenable Attack Surface Management ⓘ

changelog · Jul 16, 2026 BLOG ↗

Tenable Attack Surface Management discovers and monitors an organization's internet-facing assets, identifying exposures and vulnerabilities across the external attack surface.

Tenable ASM adds a built-in subscription to identify external-facing Microsoft SharePoint Server assets.

└──▷ HOW TO FIND IT

Identify all external-facing SharePoint Server instances in your attack surface to assess exposure to SharePoint-targeting CVEs.

> 📍 In the Tenable ASM console, navigate to Subscriptions and enable the built-in subscription labeled 'Microsoft Sharepoint Server - v1'.

› Adds built-in subscription 'Microsoft Sharepoint Server - v1' to detect and inventory external-facing SharePoint Server assets.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta ⓘ snapshot-20260717 NOTES ↗

Immuta is a data security platform that automates access control, privacy protection, and policy enforcement across cloud data environments.

Immuta adds managed groups for system-wide policies, automated Redshift discovery, OAuth 2.0 for Alation, and Delta Sharing row/column enforcement.

└──▷ HOW TO FIND IT

Write a subscription policy that allows every registered identity (agents and users) to subscribe to non-sensitive data sources, without maintaining manual group membership.

📍 In the Immuta UI, go to Policies › Subscription Policies › New Policy, set the subject to the built-in group 'All Identities', and configure the access condition for your non-sensitive data sources.

- › Automated sensitive data discovery now available for Immuta's Amazon Redshift viewless integration, supporting regex, dictionary, and column-name-regex identifiers.
- › Row filtering and column masking now enforced for Databricks Delta Sharing recipients on Unity Catalog managed, external, and foreign tables.
- › Manual data policy resync in Snowflake data source health checks now resyncs both data policies and subscription policies.
- › Integration error banners replaced by error visibility through the new connections experience.
- › Snowflake back-end architecture update improves efficiency and performance at higher scale for the Snowflake integration.

+2 more

└──▷ BREAKING ON UPGRADE

- ! Immuta policy role names used to grant users access to Snowflake data sources will be renamed as part of the Snowflake architecture update beginning July 23, 2026; existing policy roles will be recreated under new names.
- ! When a data source is disabled in Immuta, Immuta-managed grants for data owners in the data platform will be revoked (behavior change release: on by default 8/1, enabled for all 9/1).

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Ping Identity ⓘ 16 Jul 2026 (2026-07-17) BLOG ↗

Ping Identity is an identity security platform that provides authentication, single sign-on, multi-factor authentication, and access management solutions to protect users and resources.

Ping Identity AIC adds optional custom report parameters, ordered promotions API results, and multi-line journey annotations.

- Supports marking parameters in custom reports as optional, enabling more flexible report execution.
- Promotions API now returns report metadata ordered by creation date, most recent first.
- Journey annotation notes now render multi-line formatting in view mode, consistent with edit mode.

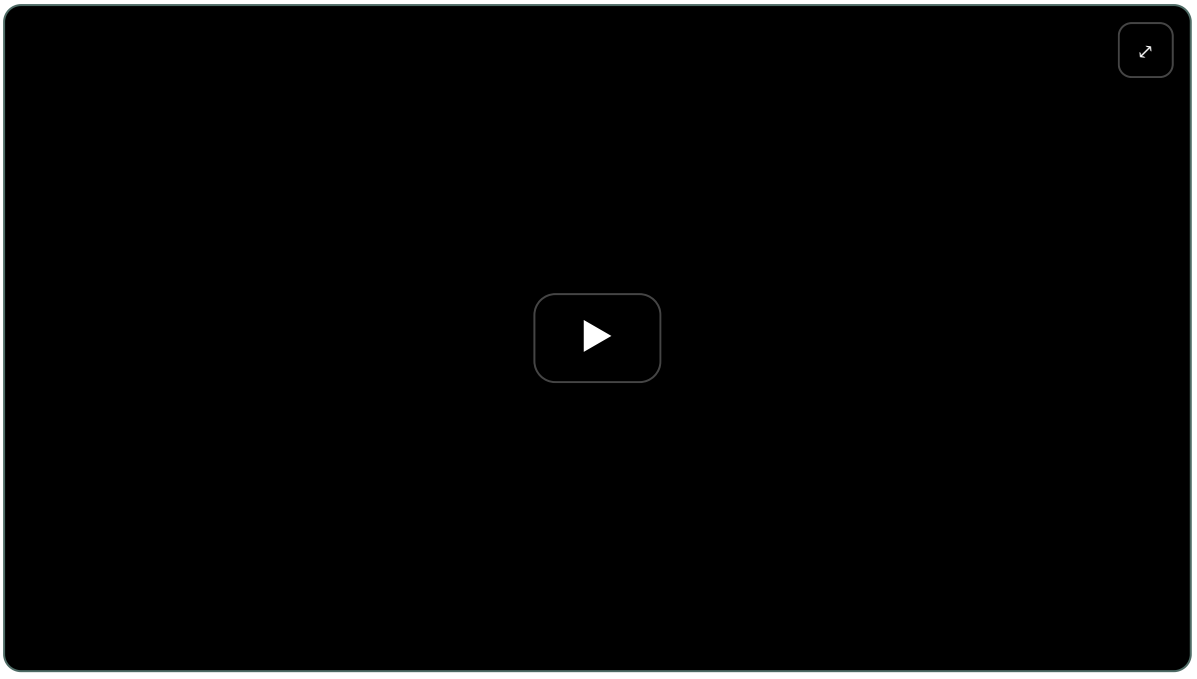
WAS THIS USEFUL? 👍 👎

Keycloak ⓘ

Multi-Cluster v2 and Stateless Mode now in Preview (2026-07-17) BLOG ↗

Keycloak is an open-source identity and access management solution that provides single sign-on, user authentication, and authorization for applications and services.

Keycloak 26.7 adds stateless mode preview, moving auth sessions and brute-force counters to the DB to enable simpler multi-cluster deployments.



└──▷ **TRY IT**

Start a single Keycloak node (or cluster) with stateless mode enabled to persist auth sessions and brute-force counters in PostgreSQL.

```
$ bin/kc.sh start \
  --features=stateless \
  --spi-cache-embedded--default--cluster-name=CLUSTER1 \
  --db=postgres --db-url=jdbc:postgresql://db-host/keycloak \
  --db-username=keycloak --db-password=<password>
```

Add a second independent cluster pointing at the same synchronously replicated database to achieve multi-cluster v2 failover with no Infinispan.

```
$ bin/kc.sh start \
  --features=stateless \
  --spi-cache-embedded--default--cluster-name=CLUSTER2 \
  --db=postgres --db-url=jdbc:postgresql://db-host/keycloak \
  --db-username=keycloak --db-password=<password>
```

- New `stateless` preview feature stores authentication sessions, action tokens, and brute-force counters in the database instead of distributed Infinispan caches.
- Enables Multi-Cluster v2: two or more independent Keycloak clusters connected only through a synchronously replicated database, with no external Infinispan cluster required.
- Cross-cluster cache invalidation now uses a database outbox/queuing table with a default polling interval of 100 ms, eliminating direct network connections between clusters.
- Authentication sessions survive full cluster restarts during minor upgrades, keeping users mid-login from losing progress.
- Multi-cluster failover is now environment-agnostic — not tied to Kubernetes or a specific cloud provider — using any synchronously replicated database (tested with Amazon Aurora PostgreSQL) and an external load balancer (tested with AWS Global Accelerator).

+1 more

WAS THIS USEFUL? 👍 👎

BeyondTrust

Base 8.3.0 release notes (2026-07-16)

2 RELEASES · 2026-07-16 | BLOG ↗ | LATEST | STABLE

BeyondTrust is a privileged access management platform that secures, controls, and monitors privileged accounts, credentials, and remote access across enterprise environments.

BeyondTrust Base 8.3.0 brings OCSF-structured syslog, a unified patch inventory view, and upload progress visibility.

- Remote syslog output now follows OCSF structure, enabling richer, standardized event forwarding to SIEM platforms with new log categories for on-premises appliances.
- Local file uploads now show an estimated progress bar instead of a generic spinner, giving clearer visibility into upload duration.
- Removes the opt-in default certificate mechanism, simplifying certificate configuration.

└──▷ BREAKING ON UPGRADE

- ! The ability to opt out of seamless (no-downtime) critical security updates is removed; these updates now apply automatically on on-premises appliances.
- ! If an appliance has not successfully communicated with the update server within a defined time window, key functionality is degraded or disabled.

▾ 1 more release in this issue · 2026-07-16

Remote Support 26.2.1 release notes (2026-07-16) | BLOG ↗ | STABLE

BeyondTrust Remote Support 26.2.1 adds multi-group Jump Client installs, FIDO2 Group Policy enforcement, AI session summaries, syslog streaming, and ARM-native support.

└──▷ TRY IT

Deploy a Jump Client into multiple Jump Groups simultaneously at install time, avoiding post-deployment API calls to reassign group membership.

```
$ msiexec /i JumpClient.msi JC_JUMP_GROUP="group-id-1,group-id-2,group-id-3" /quiet
```

- Assigns a Jump Client to multiple Jump Groups at installation time via a comma-separated `JC_JUMP_GROUP` installer parameter, eliminating post-deployment API calls to reassign group membership.
- Enforces FIDO2 phishing-resistant MFA at the Group Policy level, using the same model as existing TOTP enforcement; FIDO2 is also now supported on Pathfinder

deployments.

- Adds search, filtering, CSV export, and API access to the Duplicate Jump Clients report for programmatic duplicate detection and resolution at scale.
- Automatically generates AI-powered session summaries on session end, including case summary, step-by-step resolution, system tags, and key highlights; summaries are editable, exportable in multiple formats, multi-language, and integrate with Salesforce and ServiceNow.
- Brings vault discovery and managed/shadow account capabilities natively into the product, including a filter dropdown on the Vault > Endpoints page, new import and auto-create options for SSH, RDP, and Password Safe Jump Items, and a direct-access Password Safe Discovery method.

+8 more

BREAKING ON UPGRADE

- The ability to opt out of seamless (no-downtime) critical security updates is removed; these updates are now applied automatically on upgrade.
- RADIUS authentication provider is deprecated: existing configurations can be edited but new ones cannot be created — migrate to SAML or LDAP before future versions remove the provider entirely.

WAS THIS USEFUL? 👍 👎

SailPoint

Closed Preview: Next-Generation Mainframe Connectors for RACF, ACF2, and Top Secret (2026-07-16)

BLOG

SailPoint is an identity security platform that manages and governs user access, entitlements, and permissions across enterprise applications and systems.

SailPoint previews next-gen mainframe connectors for RACF, ACF2, and Top Secret that eliminate the Connector Gateway layer.

- Introduces direct connectivity from SailPoint Identity Management products to the SailPoint agent on the mainframe, removing the requirement to deploy and maintain a Connector Gateway.
- Adds new security and maintainability improvements to the RACF, ACF2, and Top Secret integrations on z/OS.
- Includes migration tooling to help transition existing configurations and customizations from the legacy connector to the new architecture.

WAS THIS USEFUL?

Okta Identity Engine

Weekly release 2026.07.2 (2026-07-15)

BLOG

Okta Identity Engine is a platform that enables developers to build flexible, customizable authentication and identity management workflows into their applications.

Okta adds cross-app access for AI agents, GA agent-to-agent connections, and new device OS account endpoints.

- Enables cross-app access for AI agents and applications.
- Agent-to-agent connections reaches General Availability in Production.
- New endpoints for device OS accounts now available in Early Access (Preview).

WAS THIS USEFUL?

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale ⓘ CLI 6.37.39 (2026-07-16) BLOG ↗

RegScale is a continuous compliance automation platform that helps organizations manage security controls, audits, and regulatory requirements in real time.

RegScale 6.37.39 adds eMASS asset inventory sync and expanded stakeholder address fields in the API.

› New eMASS API hardware and software baseline sync commands pull asset inventory from eMASS, create or update RegScale assets, and push new assets back to eMASS.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro ⓘ Introspect Subagent and Global Hooks NOTES ↗

Kiro is an AI-powered IDE by AWS that uses "specs" to turn requirements into code through structured, agentic software development workflows.

Kiro CLI v3 gains an Introspect subagent for in-line feature guidance and global hooks that fire across all workspaces.

└──▷ TRY IT

Ask the Introspect subagent how to configure a custom hook without leaving your terminal.

```
$ kiro-cli --v3
```

- New Introspect subagent answers questions about Kiro's features, explains custom agents, hooks, and steering files, and suggests workflow configurations — all inline (v3 only).
- Global hooks defined in `~/.kiro/hooks/` now automatically apply to every workspace, eliminating per-project duplication for cross-cutting behaviors like linting on save or pre-commit security checks (v3 only).

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ rust-v0.145.0-alpha.20

3 RELEASES · 2026-07-16 → 2026-07-17 NOTES ↗ CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Amazon Bedrock provider now supports custom transports: override base_url, auth, and HTTP headers alongside AWS profile and region.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.145.0-alpha.20
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.20
```

- Enables custom `base_url`, `auth`, and `http_headers` overrides on the built-in `amazon-bedrock` provider for custom transport scenarios.
- Supports command-based bearer authentication and configured endpoints for Bedrock without requiring AWS request signing.
- Replaces the Bedrock account `credentialSource` enum with the `usesCodexManagedCredentials` boolean for consistent reporting of managed and externally authenticated configurations.

└──▷ BREAKING ON UPGRADE

! The Bedrock account `credentialSource` enum is replaced by the `usesCodexManagedCredentials` boolean; any configuration or code referencing `credentialSource` will break.

▾ 2 more releases in this issue · 2026-07-16 → 2026-07-17

rust-v0.145.0-alpha.19 NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI alpha.19 adds concurrent terminal session writes and an app-server metadata API

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.145.0-alpha.19
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.19
```

- Enables parallel `write_stdin` tool calls across independent terminal sessions, with per-session serialization to keep output buffers consistent.
- Adds experimental `app/read` API for fetching metadata for up to 100 app IDs, with optional tool summaries, deduplication, caching, and partial-miss reporting.
- Validates reasoning effort against model metadata after agent roles override model settings, catching unsupported combinations at spawn time.
- Stabilizes the active-turn environment so in-progress turns are unaffected by settings updates during deferred execution.
- Requires `data:` scheme URLs for code-mode image output, rejecting remote HTTP URLs and other malformed image sources.

`rust-v0.145.0-alpha.18` NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI alpha.18 unifies multi-agent config under `agents`, adds encrypted MCP content support, and improves skill-selection with BM25, n-gram, and multi-query methods.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.145.0-alpha.18
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.18
```

- Adds `agents.enabled` config key and unifies multi-agent settings (concurrent thread limit, subagent model, reasoning effort) under a single `[agents]` section, with resolved values visible in TUI debug-config output.
- Adds `agents.max_concurrent_threads_per_session` (with `agents.max_threads` as an alias) to control the spawned-thread limit across both multi-agent backends.
- Adds `agent_default_subagent_model` and `agent_default_subagent_reasoning_effort` config fields so spawned agents honor configured model defaults.

> - Exposes the `agent_type` override in the `spawn_agent` schema only when agent roles are configured, avoiding an inapplicable argument in the model-visible tool.
> - Preserves encrypted MCP tool output: text content marked with `_meta["codex/encryptedContent"] = true` is now converted to `encrypted_content` function-call output items.
>
> +4 more

WAS THIS USEFUL? 👍 👎

Cline ⓘ sdk/sdk/v0.0.63 2 RELEASES · 2026-07-16 → 2026-07-17 NOTES ↗

LATEST STABLE

Autonomous coding agent as an SDK, IDE extension, or CLI assistant.

Cline SDK v0.0.63 emits `task.mistake_limit_reached` telemetry across all hosts when the consecutive-mistake limit is hit.

└──▷ GET THIS VERSION

```
$ git clone --branch sdk/sdk/v0.0.63
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout sdk/sdk/v0.0.63
```

› Adds `task.mistake_limit_reached` telemetry event emitted by the session runtime whenever the consecutive-mistake limit is reached, captured uniformly across CLI, VS Code extension, and hub daemon — including unattended auto-stops.

▾ 1 more release in this issue · 2026-07-16 → 2026-07-17

v4.0.9 NOTES ↗ CODE ↗ STABLE

Cline v4.0.9 adds GPT-5.6 ChatGPT subscription models to the provider lineup.

└──▷ GET THIS VERSION

```
$ git clone --branch v4.0.9 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout v4.0.9
```

› Adds GPT-5.6 ChatGPT subscription models as selectable options in the OpenAI provider.

WAS THIS USEFUL? 👍 👎

Claude Code ⓘ v2.1.212 NOTES ↗ CODE ↗

Claude Code is an agentic coding tool that lives in your terminal, understands your codebase, and helps you code faster by executing routine tasks, explaining complex code, and handling git workflows - all through natural language commands.

Claude Code v2.1.212 adds background `/fork`, auto-mode reset, per-session WebSearch/subagent caps, and auto-backgrounding slow MCP calls.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.1.212 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.212
```

└──▷ TRY IT

Reset auto-mode to its defaults non-interactively in a script or CI context.

```
$ claude auto-mode reset --yes
```

Cap WebSearch calls and subagent spawns to tighter limits for a long-running autonomous session to prevent runaway loops.

```
$ CLAUDE_CODE_MAX_WEB_SEARCHES_PER_SESSION=50 CLAUDE_CODE_MAX_SUBAGENTS_PER_SESSION=20 claude
```

Lower the MCP auto-background threshold to 30 seconds so slow MCP calls yield the session sooner.

```
$ CLAUDE_CODE_MCP_AUTO_BACKGROUND_MS=30000 claude
```

- Redesigns `/fork` to copy the current conversation into a new background session (visible in `claude agents`) rather than an in-session subagent; the old in-session subagent behavior is now `/subtask`.
- Adds `claude auto-mode reset` command to restore default auto-mode configuration, with an optional `--yes` flag to skip the confirmation prompt.
- Adds a per-session WebSearch call cap (default 200, tunable via

`CLAUDE_CODE_MAX_WEB_SEARCHES_PER_SESSION`) to prevent runaway search loops.

› Adds a per-session subagent spawn cap (default 200, override via `CLAUDE_CODE_MAX_SUBAGENTS_PER_SESSION`); `/clear` resets the budget.

› MCP tool calls exceeding 2 minutes automatically move to the background to keep the session responsive; threshold is configurable or disableable via `CLAUDE_CODE_MCP_AUTO_BACKGROUND_MS`.

+8 more

└──▷ BREAKING ON UPGRADE

! The Task tool's `mode` parameter is now ignored; subagents inherit the parent session's permission mode by default.

WAS THIS USEFUL? 👍 👎

Pi ⓘ v0.80.9 NOTES ↗

AI agent toolkit: unified LLM API, agent loop, TUI, coding agent CLI

Pi v0.80.9 adds Kimi K3 support with deferred tool loading across five providers and switches xAI default to Grok 4.5.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.80.9 https://github.com/earendil-works/pi.git
# already have the repo? check out this version:
$ git checkout v0.80.9
```

- Adds Kimi K3 model support across Kimi Coding, Moonshot AI, Moonshot AI China, OpenRouter, and Vercel AI Gateway.
- Adds Kimi deferred tool loading for progressive, extension-driven tool activation via Kimi's native protocol.
- Changes the default xAI model to Grok 4.5 and updates xAI login to use a prefilled device-authorization link labeled 'Sign in with SuperGrok or X Premium'.

└──▷ BREAKING ON UPGRADE

! Grok 3, Grok 3 Fast, Grok 4.20 variants, and Grok Code Fast 1 have been removed from the built-in xAI model catalog — any configuration referencing these models will break.

WAS THIS USEFUL? 👍 👎

OpenCode ⓘ v1.18.3 NOTES ↗ CODE ↗

The open source coding agent.

OpenCode v1.18.3 adds Up Arrow to close the subagent picker and lets the command palette search sessions from the home page.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.18.3
https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.18.3
```

- Adds Up Arrow shortcut to close the subagent picker when the first item is selected, enabling faster keyboard-only navigation.
- Enables the home page command palette (Cmd+K) to find and open existing sessions, filtered by the currently selected server.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

CrewAI ⓘ 1.15.4 3 RELEASES · 2026-07-16 → 2026-07-17 NOTES ↗ LATEST

STABLE

Framework for orchestrating role-playing, autonomous AI agents. By fostering collaborative intelligence, CrewAI empowers agents to work together seamlessly, tackling complex tasks.

CrewAI 1.15.4 promotes the Skills Repository to stable, making AI coding-agent guidance an officially supported feature.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.15.4 https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.4
```

└──▷ TRY IT

Install the now-stable CrewAI Skills into Claude Code so it automatically applies CrewAI best practices when you scaffold Flows, configure Crews, or design agents.

```
$ /plugin marketplace add crewAIInc/skills
/plugin install crewai-skills@crewai-plugins
/reload-plugins
```

Install the stable CrewAI Skills into Cursor, Codex, Windsurf, or any other coding agent via the skills.sh registry.

```
$ npx skills add crewaiinc/skills
```

› Promotes the Skills Repository out of experimental status, making structured coding-agent instructions (for Claude Code, Cursor, Windsurf, etc.) a stable, supported capability.

▾ 2 more releases in this issue · 2026-07-16 → 2026-07-17

> 1.15.3 NOTES ↗ STABLE
>
> CrewAI 1.15.3 adds step interception hooks with @on, org ID support in PlusAPI, and declarative Flow execution in the TUI.
>
> └──▷ GET THIS VERSION

```
$ git clone --branch 1.15.3 https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.3
```

└──▷ **USE IT**

Intercept execution boundaries in a crew or flow to log, audit, or gate agent activity at specific lifecycle points.

```python
from crewai import on

@on("before_llm_call")
def audit_llm_input(context):
    print("LLM about to be called with:", context)

@on("after_llm_call")
def log_llm_output(context):
    print("LLM responded with:", context)
```

- Adds step interception points and a generic hook dispatcher, enabling fine-grained execution-boundary callbacks via `@on` decorators.
- Adds `organization_id` parameter to the PlusAPI client for multi-org deployments.
- Supports running declarative Flows on the TUI with headless terminal fallback.

1.15.3a1 NOTES ↗ PRE-RELEASE

CrewAI 1.15.3a1 adds step interception hooks via `@on`, declarative flow support in the TUI, and an organization ID parameter for the PlusAPI client.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 1.15.3a1
https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.3a1
```

- Adds step interception points and a generic hook dispatcher, enabling `@on - decorated` callbacks at execution boundaries for fine-grained observability and control.
- Supports running declarative flows on the TUI with a headless terminal fallback.

- Adds an organization ID parameter to the PlusAPI client for multi-org deployments.
- Improves custom OpenAI URL handling for self-hosted or proxy LLM endpoints.

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

Ollama ⓘ v0.32.1 NOTES ↗ CODE ↗

Get up and running with Kimi-K2.6, GLM-5.1, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.

Ollama v0.32.1 adds CWD-aware agent context, actionable web-auth guidance, and MLX load-timeout support.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.32.1 https://github.com/ollama/ollama.git
# already have the repo? check out this version:
$ git checkout v0.32.1
```

- Adds current working directory to the interactive agent's system prompt for better project context.
- Agent web search and fetch now surface an actionable `ollama signin` prompt when authentication is required, instead of a generic error.
- MLX text model loading now respects the `OLLAMA_LOAD_TIMEOUT` environment variable.
- Improves Gemma 4 tool calling and multi-turn reasoning, including more reliable tool-response continuations.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe