THE CYBER TOOLCHAIN

NO. 017

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 16, 2026 · EVERY DAY

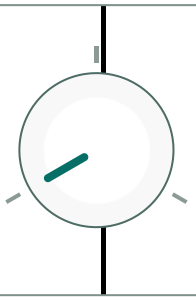
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 37 tools

APPLICATION SECURITY

vim ? help

◆ SAST / SCA / Code Security

StackHawk

Feedback ⓘ snapshot-20260715 NOTES ↗

StackHawk is a dynamic application security testing (DAST) tool that helps developers find and fix security vulnerabilities in APIs and web applications during development.

StackHawk adds self-upgrade, scan drill-down commands, standalone binary, AI skills status, and org-wide Findings page.

└──▷ TRY IT

Check whether a newer HawkScan version is available before committing to an upgrade in a CI pre-flight step.

```
$ hawk upgrade --check
```

Iterate on an auth script against a live staging URL without running a full scan, with auto-reload on save.

```
$ hawk validate script --watch https://staging.example.com
```

Inspect the exact attack surface (including GraphQL operations) exercised by a completed scan to verify coverage.

```
$ hawk op scan uris --app my-api --graphql
```

New `hawk upgrade` command updates HawkScan CLI in place, detecting install method (Homebrew, binary, MSI, PKG, Docker) and printing the exact command when direct upgrade isn't possible; `--check` previews available versions and `--version <x.y.z>` pins to a specific release.

New `hawk skills status` command reports which StackHawk AI agent skills (Claude Code, Codex, Copilot) are installed and whether they are current, fetching latest published versions live.

New `hawk op scan config` retrieves the exact effective configuration a completed scan ran with; `hawk op scan uris` lists the full attack surface including GraphQL operations and JSON-RPC methods.

`--app` flag on all `hawk op` commands now accepts either an application name or a UUID, eliminating the need to look up UUIDs separately.

New `hawk validate script` command runs HawkScript auth logic against a single target URL without a full scan; `--watch` enables daemon mode that re-executes on every file save.

+12 more

└─▷ **BREAKING ON UPGRADE**

- ! The `hawk-{version}.zip` download format is removed; pipelines downloading the ZIP must migrate to the OS/architecture-specific binary at `https://download.stackhawk.com/hawk/{VERSION}/{PLATFORM}/hawk` (Windows: `hawk.exe`).
- ! StackHawk CircleCI Orb v2.0.0: job and parameter names are renamed to snake_case — e.g., `stackhawk/hawkscan-local` becomes `stackhawk/hawkscan_local` and `app-id` becomes `app_id`. Pipelines pinned to a 1.x version are unaffected until upgraded.

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ

New Control: Jobs Should Use a Secure Registry Instead of Public Package Registries (202

BLOG ↗

StepSecurity is a platform that secures CI/CD pipelines by detecting and preventing supply chain attacks, hardening GitHub Actions workflows, and monitoring runtime behaviors.

StepSecurity adds a secure-registry CI/CD control and OIDC replaceability tracking for Actions secrets.

└──▷ **HOW TO FIND IT**

Identify CI/CD jobs pulling packages from public registries so you can route them through a secure proxy before a supply-chain compromise occurs.

> 📍 In the StepSecurity console, go to Overview dashboard › Controls, open the `Jobs Should Use a Secure Registry Instead of Public Package Registries` control, then filter findings by repository or workflow to prioritize remediation.

Find secrets that can be eliminated entirely by switching to OIDC, reducing long-lived credential exposure across your repositories.

> 📍 In the StepSecurity console, go to Actions Secret › Repository Secrets tab, then enable the `OIDC Replaceable` filter to list only secrets (e.g., `NPM_TOKEN`) that trusted publishing can replace.

New Overview dashboard control detects CI/CD jobs that download packages directly from public registries (`registry.npmjs.org` , pypi.org, registry-1.docker.io) instead of a secure proxy registry.

Control findings are filterable by repository, workflow, and compliance status, with remediation tracking and finding suppression on the control detail page.

Repository Secrets tab now surfaces Usage Status (Active, Stale, Unused, Unknown) for each secret, including which workflow file and step references it and the last run time and conclusion.

Secrets that can be replaced with OIDC (e.g., `NPM_TOKEN` replaceable via npm trusted publishing) are flagged in the Actions Secret section.

New summary cards on the Repository Secrets tab show Total Secrets, Unused Secrets, Stale Secrets, and OIDC Replaceable counts for an at-a-glance secrets posture view.

+1 more

WAS THIS USEFUL?  

◆ Web Application Security

Cloudflare Feedback ⓘ

A broken DNSSEC rollover took down .al. Now 1.1.1.1 tells you when validation is bypasse

BLOG ↗

Cloudflare is a cloud platform that provides website security, performance optimization, and DDoS protection by routing traffic through its global network.

Returns EDE 9 (DNSKEY Missing) alongside EDE 33 to expose the underlying DNSSEC chain-of-trust failure that triggered the NTA.

Emits EDE 33 on every response under an active NTA — including non-DNSSEC domains — so the entire affected zone is transparently flagged.

WAS THIS USEFUL?  

PortSwigger Burp Suite Feedback ⓘ

Professional / Community 2026.7 (2026-07-16) BLOG ↗

Burp Suite is a web security testing platform that intercepts, analyzes, and manipulates HTTP/S traffic to identify vulnerabilities in web applications.

Burp Suite 2026.7 launches beta of AI-driven agentic pentesting that autonomously pursues user-defined testing goals.

WAS THIS USEFUL?  

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft Feedback ⓘ

Elcomsoft Phone Breaker 11.03: universal binary for Macs, SMS two-factor authentication

BLOG ↗

ElcomSoft develops forensic and password recovery tools that help law enforcement, investigators, and security professionals access encrypted files, devices, and accounts.

Elcomsoft Phone Breaker 11.03 adds universal binary for macOS and restores SMS-based iCloud 2FA sign-in.

Ships as a universal binary, running natively on both Apple Silicon and Intel Macs without Rosetta translation.

Removes the pre-login 2FA status check that Apple no longer supports, enabling successful authentication flows.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Palo Alto Networks

Feedback ⓘ changelog · Jul 16, 2026

BLOG ↗

Palo Alto Networks is a cybersecurity platform that provides advanced firewalls, cloud security, and AI-driven threat detection to protect networks and enterprises.

Prisma AIRS AI Gateway reaches GA with inline LLM/MCP/A2A enforcement, agent identity binding, and quantum-resilient SASE fabric via AT&T partnership.

.

Introduces Prisma AIRS AI Gateway as a unified control plane sitting inline between AI apps, LLMs, MCP tool calls, and A2A interactions for centralized governance.

Adds unified observability across all AI usage: tracking users, projects, models, token counts, latency, and cost in a single view to expose shadow AI.

Enables central governance of approved models, tools, and access with FinOps controls including per-team/project budget enforcement and rate limiting before access is granted.

Adds Coding Assistant Security that replaces raw provider API keys with scoped, per-user/team credentials to protect secrets and proprietary code.

Introduces Agent Identity Security that binds verifiable, ephemeral identities to agents at execution time, enforcing just-in-time least-privilege access.

+7 more

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.34.0 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 5.34.0 adds E2E Networks provider, 9 new security checks across AWS/Azure/Kubernetes/STACKIT, and cross-provider compliance rollup.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 5.34.0 https://github.com/prowler-cloud/prowler.git
# already have the repo? check out this version:
$ git checkout 5.34.0
```

└──▷ TRY IT

Scan your E2E Networks environment for misconfigurations across compute, networking, and storage after setting your API credentials.

```
$ export E2E_NETWORKS_API_KEY="your-api-key"
export E2E_NETWORKS_AUTH_TOKEN="your-auth-token"
export E2E_NETWORKS_PROJECT_ID="your-project-id"
prowler e2enetworks
```

New provider: Scans E2E Networks with 27 checks covering compute nodes, networking, security groups, load balancers, block/file storage, and managed databases.

New AWS check `ec2_ami_account_block_public_access`: verifies AMI block public access is enabled at the account level per region.

New AWS check `datapipeline_pipeline_no_secrets_in_definition`: scans Data Pipeline object fields, parameter objects, and parameter values for hardcoded secrets using Kingfisher.

New AWS check `elbv2_listener_pqc_tls_enabled`: verifies ELBv2 HTTPS/TLS listeners use post-quantum TLS security policies (TLS 1.2+) to reduce harvest-now-decrypt-later exposure.

New AWS check `amplify_app_no_secrets_in_environment`: scans Amplify app and branch environment variables and build settings (buildSpec) for hardcoded secrets

using Kingfisher.

+6 more

└─▷ **BREAKING ON UPGRADE**

! UI integration enable flags renamed: `UI_SENTRY_ENABLE` → `UI_SENTRY_ENABLED`, `UI_GOOGLE_TAG_MANAGER_ENABLE` → `UI_GOOGLE_TAG_MANAGER_ENABLED`, `UI_POSTHOG_ENABLE` → `UI_POSTHOG_ENABLED`; deployments that set the former names must update them.

WAS THIS USEFUL?  

Sysdig Secure

Feedback ⓘ Secure SaaS - July 16, 2026 BLOG ↗

Sysdig Secure is a cloud-native security platform that detects threats, manages vulnerabilities, ensures compliance, and protects containers, Kubernetes, and cloud environments at runtime.

Sysdig Secure now supports multiple S3/S3-compatible custom storage buckets across Captures, Response Actions, and Rapid Response.

└──▷ **HOW TO FIND IT**

Connect a MinIO bucket to use as storage for Captures and Rapid Response sessions.

> 📍 In the console, go to Integrations › Data Sources › Storage, click Add Bucket, select S3-compatible endpoint, enter your MinIO endpoint URL and credentials, then click Test Connection and Save.

Supports connecting multiple S3 or S3-compatible buckets (MinIO, Google Cloud Storage, and others) for use across Captures, Response Actions, and Rapid Response.

New Custom Storage configuration supports AWS access keys, AWS role delegation (AssumeRole), and S3-compatible endpoints, with connection testing on save and per-bucket failure history.

└──▷ **BREAKING ON UPGRADE**

! The previous single shared S3 Capture Storage configuration is replaced by Custom Storage; existing setups must be migrated to the new multi-bucket configuration under Integrations > Data Sources > Storage.

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

NVIDIA OpenShell

Feedback ⓘ v0.0.85 NOTES ↗

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

OpenShell v0.0.85 adds elevated gateway info surfacing to the gateway component.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.0.85 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.85
```

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Anchore syft

Feedback ⓘ v1.48.0 NOTES ↗

CLI tool and library for generating a Software Bill of Materials from container images and filesystems

Syft v1.48.0 adds vcpkg, macOS .app, Kerberos 5, SafeTensor, and Ubuntu Pro/ESM cataloging plus Go binary symbol capture.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.48.0 https://github.com/anchore/syft.git
# already have the repo? check out this version:
$ git checkout v1.48.0
```

Adds dependency relationship edges from `mix.lock` to SBOM output.

Adds SafeTensor model type cataloging to SBOM output for ML model inventories.

Captures symbols from Go binaries for richer binary-level SBOM data.

Detects Ubuntu Pro/ESM extended support status on scanned systems.

Adds a scan duration timer to Syft output.

+4 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Exabeam

Feedback ⓘ Enhanced Threat Center Endpoint (2026-07-15)

BLOG ↗

Exabeam is a security information and event management (SIEM) platform that uses AI and behavioral analytics to detect, investigate, and respond to cybersecurity threats.

Exabeam Threat Center adds `caseClosedTimestamp` field to case details, enabling Time-to-Close metrics.

Adds `caseClosedTimestamp` (date-time string) to the Get Case Details response schema, exposing when a case was closed for Time to Close (TTC) and Mean Time to Close (MTTC) calculations.

WAS THIS USEFUL? 👍 👎

Google SecOps

Feedback ⓘ July 15, 2026 (2026-07-15) BLOG ↗

Google SecOps is a cloud-based security operations platform that helps organizations detect, investigate, and respond to cybersecurity threats using AI-powered analytics and threat intelligence.

Google SecOps dashboards gain Advanced Filtering with dynamic query variables, regex support, and multi-select logic injection into YARA-L at runtime.

└──▷ HOW TO FIND IT

When investigating across multiple threat actor IPs, use a multi-select advanced filter token to fan out a single YARA-L dashboard query across all selected values without editing the query manually.

> 📍 In the dashboard editor, go to Filters › Add Advanced Filter, define a Token Variable (e.g., `source_ip`), set Filter Value Generation to a YARA-L query result or static list, enable Multi-Select, and set the delimiter to `|` — then reference the token in your YARA-L rule where the value should be injected at runtime.

Adds Advanced Filtering in dashboards (Public Preview): security analysts can define token variables to inject dynamic values, complex regular expressions, or boolean logic directly into YARA-L queries at runtime.

Supports token values generated dynamically from YARA-L query results or entered as a static manual list.

Enables customizable prefix/suffix wrappers around token values to express regex or other query logic.

Supports multi-select tokens with a configurable delimiter (e.g., `|`) for combining multiple values in a single YARA-L query.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Fleet

Feedback ⓘ fleet-v4.89.0 NOTES ↗

Open device management

Fleet 4.89.0 adds policy label targeting, setup experience policy checks, new vulnerability filtering, Google Workspace IdP integration, and more.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch fleet-v4.89.0
https://github.com/fleetdm/fleet.git
# already have the repo? check out this version:
$ git checkout fleet-v4.89.0
```

└──▷ TRY IT

Filter vulnerable software by CVSS score and known exploit status on a device's self-service software tab — useful for prioritizing critical, actively exploited vulnerabilities.

```
$ GET /device/{token}/software?min_cvss_score=7.0&exploit=true
```

Reset a policy's pass/fail results and membership immediately — useful after remediating a widespread misconfiguration before the next scheduled check.

```
$ POST /api/v1/fleet/policies/42/reset
```

Adds ability to target policies to hosts using a combination of `include` and `exclude` labels.

Adds policy check before installing Windows and Linux setup experience software, skipping installs when the policy already passes to speed up end-user setup.

Changes calendar remediation events to be scheduled on the next business day (skipping weekends) after a policy failure.

Adds automation runs table to the policy details page showing per-host automation outcomes with filtering, search, and a reset policy action.

Adds per-host activity log entries when policy automations (webhook, tickets, Google Calendar, and Microsoft conditional access) fail or succeed.

+23 more

└──▷ **BREAKING ON UPGRADE**

! RFC 1918 and other private network ranges are now blocked by default in Fleet's HTTP client; deployments with on-prem integrations (e.g. EJBCA, Jira, SCEP servers on private networks) must add `--allow_private_network_integrations` to restore connectivity.

WAS THIS USEFUL? 👍 👎

Microsoft Defender

Feedback ⓘ changelog · Jul 15, 2026

BLOG ↗

Microsoft Defender is a built-in Windows security tool that protects devices from viruses, malware, ransomware, and other threats through real-time monitoring and scanning.

Microsoft launches Defender Experts Threat Intelligence service and expands MDR to multi-cloud/third-party coverage via Sentinel.

New Microsoft Defender Experts Threat Intelligence service delivers curated, expert-interpreted threat intelligence tailored to your industry, geography, and environment — including early-warning campaign alerts, evolution updates, and recurring expert briefings.

Microsoft Defender Threat Intelligence (MDTI) capabilities are now fully converged into the Defender portal, making intelligence available in real-time across detection, investigation, response, hunting, and automation without context switching.

Microsoft Defender Experts MDR (formerly Defender Experts for XDR) adds a new Plan 2 tier with third-party and multi-cloud coverage powered by Microsoft Sentinel, enabling expert-led monitoring and response across non-Microsoft cloud, identity, email, network, and endpoint sources.

Defender Experts MDR Plan 2 provides cross-platform threat analysis correlating Microsoft and non-Microsoft signals into a single incident narrative with vendor-aware guidance.

└──▷ **BREAKING ON UPGRADE**

! Microsoft Defender Experts for XDR is renamed to Microsoft Defender Experts MDR (existing capabilities carry forward as Plan 1).

WAS THIS USEFUL? 👍 👎

CrowdStrike

Feedback ⓘ undated-fa98e67b BLOG ↗

CrowdStrike is a cloud-native cybersecurity platform that protects endpoints, cloud workloads, and identities from cyber threats using AI-powered threat detection and response.

CrowdStrike launches Falcon AIDR, a new runtime security category detecting and blocking threats targeting and originating from AI agents across endpoint, SaaS, and cloud.

Introduces Falcon AI Detection and Response (AIDR), a new security category providing unified runtime visibility, detection, and automated response for AI agents across endpoints, SaaS, and cloud environments.

Discovers shadow AI by illuminating employee agent and tool usage across endpoints, AI gateways, MCP servers, cloud environments, and Copilot ecosystems.

Enforces granular attribute-based access controls aligned to AI governance policy, including policy based on user and model version.

Automatically detects and blocks confidential data, PII, secrets, keys, and regulated data before exposure, with multiple redaction methods including format-preserving encryption.

Generates comprehensive audit trails of AI runtime behavior for compliance, forensics, and continuous improvement.

+1 more

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Global approved list for teams (2026-07-15) · 2 RELEASES · 2026-07-15 · BLOG ↗

LATEST · STABLE

Trend Vision One is a unified cybersecurity platform that provides extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

Trend Vision One adds Global approved list for teams, letting admins exempt specific collaboration spaces from ATP and DLP scanning.

└──▷ **HOW TO FIND IT**

Exempt a trusted internal team from ATP and DLP scanning to reduce noise and improve performance for that collaboration space.

> 📍 In the console, go to Email and Collaboration Security → Cloud Email and Collaboration Protection → Policies, then open the Global approved list for teams and add the trusted team.

New Global approved list for teams lets administrators designate trusted teams to bypass Advanced Threat Protection and Data Loss Prevention inspection while keeping full protection elsewhere.

Reduces alert noise and improves performance by scoping security scanning away from known-safe collaboration spaces.

▾ **1 more release in this issue** · 2026-07-15

Skip notification for Pass actions (2026-07-15) · BLOG ↗ · STABLE

Trend Vision One lets admins suppress notifications for Pass-action detections with a single global toggle.

└──▷ **HOW TO FIND IT**

Reduce alert fatigue by silencing Pass-action notifications across your org — useful when high-volume benign detections are flooding admin and user inboxes.

> 📍 In the console, go to Email and Collaboration Security → Cloud Email and Collaboration Protection → Policies and enable `Skip notification for Pass actions`.

Adds a global `Skip notification for Pass actions` setting that suppresses admin and end-user alerts when a detection results in no action taken on an email or file.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir Feedback ⓘ

Tenzir Node v6.7.0: Fan out and rejoin pipelines with fork_merge (2026-07-15)

2 RELEASES · 2026-07-15 BLOG ↗ LATEST STABLE

Tenzir is an open-source security data pipeline platform that enables teams to collect, route, transform, and analyze security telemetry at scale.

Tenzir v6.7.0 adds fork_merge for fanning out subpipelines and a force-stop action for pipelines.

New `fork_merge` operator fans out subpipelines over the same input and merges their outputs into a single stream.

Adds a force-stop action for pipelines in the Tenzir Platform.

Improved Tenzir Platform performance.

Restores unstructured (raw-log) ingestion support for Google SecOps.

▾ 1 more release in this issue · 2026-07-15

> Tenzir Platform v1.37.0: Custom library sources (2026-07-15) BLOG ↗ STABLE
>
> Tenzir Platform v1.37.0 adds custom Git-backed package libraries, force-stop for pipelines, and a workspace key minting command.
>
> Adds custom library sources: organizations can add public or private Git repositories alongside the built-in Tenzir Library to maintain their own package libraries.
>
> Adds a Force stop action for pipelines that are slow to drain, enabling operators to hard-terminate stalled pipelines.
>
> Adds an admin command to mint workspace keys.
>
> └─▷ BREAKING ON UPGRADE
>
> ! The Random Example feature has been removed.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ Restrict who can create API keys (2026-07-16)

BLOG ↗

Tines is a no-code automation platform that enables security and operations teams to build powerful workflows by connecting tools and automating repetitive tasks.

Tines lets tenant owners restrict which users can create API keys via a new tenant-level setting.

└──▷ HOW TO FIND IT

Lock down API key creation to prevent non-owner users from generating credentials.

> 📍 In the Tines console, go to Settings › API keys and enable the restriction option under the tenant settings section.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

defectdojo

Feedback ⓘ 3.1.101 NOTES ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

Adds `DD_WATSON_SEARCH_ENABLED` toggle to control Watson search indexing in DefectDojo.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.1.101 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.1.101
```

└──▷ USE IT

Disable Watson search indexing to reduce database load in environments where full-text search is not needed.

```ini
DD_WATSON_SEARCH_ENABLED=False
```

New `DD_WATSON_SEARCH_ENABLED` setting in `settings.dist.py` / `local_settings.py` to gate Watson search indexing on or off.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta

Feedback ⓘ snapshot-20260716

2 RELEASES · SEEN 2026-07-15 → 2026-07-16 NOTES ↗ LATEST STABLE

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

Immuta adds automated Redshift data discovery, built-in managed groups, Delta Sharing row/column masking, Teradata hashing, and OAuth 2.0 for Alation.

└──▷ HOW TO FIND IT

Use the embedded documentation chatbot to answer governance questions without leaving the Immuta UI.

> 📍 In the Immuta UI, click the Documentation icon in the top menu, then type your question in natural language in the AI-powered documentation chatbot.

Introduces three built-in Immuta-managed groups — All Agents, All Users, and All Identities — available for use in subscription and data policies without manual membership management.

Enables row filtering and column masking (subscription and data policies) on Unity Catalog managed, external, and foreign tables in Databricks Delta Sharing shares.

Supports OAuth 2.0 authentication for Immuta's Alation integration alongside existing API key authentication, eliminating recurring token rotation.

Extends Immuta's Databricks Unity Catalog tag ingestion to foreign catalogs (Lakehouse Federation) by falling back to system.information_schema tag views.

Adds constant and hash masking support for Immuta's Teradata integration, expanding beyond existing null masking.

+4 more

> └──▷ BREAKING ON UPGRADE
>
> ! Beginning August 1, disabling a data source in Immuta will revoke data owners' Immuta-managed grants in the data platform for that data source; previously, disabling a data source did not remove those grants.
>
> ! Beginning August 1, Starburst (Trino) principal name matching becomes case-insensitive for all accounts; setups relying on case-sensitive principal name matching will be affected.

▾ 1 more release in this issue · seen 2026-07-15 → 2026-07-16

snapshot-20260715 NOTES ↗ LATEST STABLE

Immuta adds automated Redshift discovery, built-in managed groups, OAuth for Alation, Delta Sharing masking, and Teams/Slack notifications.

└──▷ **HOW TO FIND IT**

Access Immuta documentation and ask governance questions in natural language without leaving the product.

> 📍 In the Immuta UI, click the Documentation icon in the top menu to open the embedded docs or AI-powered chatbot, then type your question in natural language.

Introduces three built-in system-managed groups — All Agents, All Users, and All Identities — usable in both subscription and data policies without manual membership management.

Supports OAuth 2.0 authentication for Immuta's Alation integration, eliminating the need for recurring API key rotation.

Enables row filtering and column masking on Unity Catalog managed, external, and foreign tables in Databricks Delta Sharing shares (now GA for all accounts).

Adds Slack and Microsoft Teams notifications for the Request app (generally available), delivering event-driven alerts for approvals, data access changes, and requests.

Expands Microsoft Teams Request app notifications to the Asia Pacific region (public preview).

+4 more

> └──▷ **BREAKING ON UPGRADE**
>
> ! When a data source is disabled in Immuta, Immuta now revokes data owners' Immuta-managed grants in the data platform; previously, disabling a data source did not remove those grants (on by default 8/1, enabled for all 9/1).

WAS THIS USEFUL? 👍 👎

Zscaler Data Protection

Feedback

changelog · Jul 16, 2026

BLOG ↗

Zscaler Data Protection is a cloud-native DLP solution that prevents sensitive data leakage across web, cloud applications, email, and endpoints.

Zscaler DLP, DSPM, and inline proxy now detect and block SSLKEYLOGFILE TLS session-key material across endpoints, cloud storage, and network traffic.

Adds endpoint DLP detection for SSLKEYLOGFILE content matching legacy and newly standardized formats, including obfuscated or renamed files, to catch TLS key material exfiltration on user devices.

Extends DSPM scanning to identify files containing TLS session secrets across SaaS platforms, public cloud, and on-premises storage regardless of naming convention or file type.

Enables inline DLP engine to detect and policy-enforce outbound SSLKEYLOGFILE transfers in real time via file uploads, web channels, non-web channels, and email — with options to block, alert, or trigger investigation workflows.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Nudge Security

Feedback ⓘ

Address sprawling OAuth and browser extension risk with new Nudge Security AI agents

NOTES ↗

Nudge Security discovers and manages SaaS applications across an organization, helping security teams gain visibility and govern employee-introduced tools through automated nudges.

Nudge Security adds two AI agents to automatically assess and remediate OAuth grant and browser extension risk.

New OAuth Grant Risk Analyst AI agent reviews newly discovered OAuth grants, flags overly permissive or high-risk ones, and resolves them by automatically revoking access or nudging the grantor — with human-in-the-loop oversight.

New Browser Extension Risk Analyst AI agent inspects installed browser extension artifacts and metadata, flags malicious, compromised, or misrepresented extensions, and orchestrates disable and uninstall workflows.

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ

1Password for Linux beta 8.12.30-19 (2026-07-15)

5 RELEASES · 2026-07-14 → 2026-07-15 BLOG ↗ LATEST STABLE

1Password is a password manager that securely stores, generates, and autofills passwords, credentials, and sensitive information across devices and browsers.

1Password for Linux adds a local MCP server so AI assistants can access 1Password Developer tools.

└──▷ HOW TO FIND IT

Enable the local MCP server so an AI assistant (e.g., Claude, Cursor) can query 1Password Developer tools during a session.

> 📍 In the 1Password app, go to Settings › Developer and start the local MCP server from the Developer tab.

New local MCP server in the Developer tab lets AI assistants connect to 1Password Developer tools.

Watchtower descriptions now explicitly surface scanning for plaintext secrets, including .env files and SSH keys.

▾ 4 more releases in this issue · 2026-07-14 → 2026-07-15

1Password for Mac beta 8.12.30-19 (2026-07-15) BLOG ↗ STABLE

1Password adds a local MCP server for AI assistant integration and expands Watchtower plaintext secret scanning visibility.

└──▷ HOW TO FIND IT

Connect an AI assistant to 1Password Developer tools by starting the local MCP server from the app.

> 📍 In the 1Password app, go to Settings › Developer, then enable and start the local MCP server. Point your AI assistant's MCP client configuration to the local server endpoint exposed there.

New local MCP server in the Developer tab lets AI assistants access 1Password Developer tools.

Watchtower and Developer homepage descriptions now surface scanning coverage for plaintext secrets, including .env files and SSH keys.

1Password for Windows beta 8.12.30-19 (2026-07-15) BLOG ↗ STABLE

1Password for Windows adds a local MCP server for AI assistant access to Developer tools and expands Watchtower secret scanning descriptions.

└──▷ **HOW TO FIND IT**

Connect an AI assistant (e.g., Claude, Cursor) to your 1Password vault by starting the local MCP server from the Developer tab.

> 📍 In the 1Password app, go to Settings › Developer, then enable and start the local MCP server to expose 1Password Developer tools to your AI assistant.

Adds a local MCP server launchable from the Developer tab, enabling AI assistants to access 1Password Developer tools.

Watchtower descriptions on the Developer homepage now explicitly mention scanning for plaintext secrets, including .env files and SSH keys.

`1Password in the browser 8.12.30-19 (2026-07-15)` **BLOG ↗** **STABLE**

1Password in the browser adds toggleable extension diagnostic logging in Settings > Help.

└──▷ **HOW TO FIND IT**

Enable diagnostic logging when troubleshooting extension behavior to capture logs for 1Password Support.

> 📍 In the browser extension, go to Settings › Help › Share diagnostic logs with 1Password and toggle logging on.

Adds a toggle to enable or disable extension diagnostic logging via Settings > Help > Share diagnostic logs with 1Password.

Improves SSO account migration so the desktop app only comes to the front when the migrating account is already configured in the desktop app.

`1Password CLI Beta 2.38.0-beta.01 (2026-07-14)` **BLOG ↗** **PRE-RELEASE**

1Password CLI gains a command to manually trigger Google Workspace user provisioning syncs.

└──▷ **TRY IT**

Manually kick off a Google Workspace provisioning sync without waiting for the scheduled interval — useful after bulk directory changes.

```
$ op provisioning google sync
```

WAS THIS USEFUL? 👍 👎

iDenfy

Feedback ⓘ snapshot-20260716 NOTES ↗

iDenfy is an identity verification platform that automates KYC/AML compliance through AI-powered document scanning, facial recognition, and liveness detection for businesses.

iDenfy adds Brazil Digital CNH, new KYB automation API, Risk Assessment PDF export, Arabic/RTL IDV support, and multi-theme KYC branding.

Adds Brazil Digital CNH support for identity verification.

New API endpoint to re-run KYB automation workflows.

Dedicated tax identification fields (VAT Number, EIN, TIN) added to KYB workflows.

External Reference field now available in KYC verification reports.

Risk Assessment PDF export now supported.

+30 more

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ Choose what SAML sends as a login hint (2026-07-15)

BLOG ↗

Clerk is a complete user authentication and management platform that provides embeddable UI components, APIs, and dashboards for adding secure sign-in, sign-up, and user management to applications.

Clerk SAML connections can now send an email address, custom attribute, or no login hint to your IdP.

Supports per-SAML-connection configuration of the login hint sent to identity providers: email address, a custom attribute, or none.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Zscaler

Feedback ⓘ changelog · Jul 15, 2026 BLOG ↗

Zscaler is a cloud-native security platform that provides secure internet access, zero trust networking, and threat protection for users, devices, and applications.

Zscaler launches Autonomous Application Shield to auto-apply app-specific protections in the patch gap via App Connectors and global threat intel.

└──▷ HOW TO FIND IT

Request early access to Autonomous Application Shield to get hands-on experience and direct roadmap input before general availability.

> 📍 In the Zscaler console, contact your Zscaler representative or register via the Early Access Program webinar link to enable Autonomous Application Shield for your private applications.

New Autonomous Application Shield capability continuously assesses private application exposure via App Connectors and automatically applies per-app protections in the Zscaler cloud without requiring a patch.

Per-application protection reasoning: the Zscaler cloud evaluates each private application individually and applies only the controls that specific application needs, reducing false positives and performance impact.

Automatic protection adaptation when developers ship new releases — security policy updates without manual re-tuning, policy review meetings, or security team intervention.

Global threat intelligence integration feeds zero-day and emerging attack technique data from Zscaler's worldwide customer base into the protection engine, with ML continuously tuning policies per application.

WAS THIS USEFUL? 👍 👎

DNSFilter Feedback ⓘ

07/15 July 15 - Identities Connections Release (2026-07-15) BLOG ↗

DNSFilter is a cloud-based DNS security platform that protects users from malware, phishing, and unwanted content by filtering DNS queries in real time.

DNSFilter Identities framework adds direct Entra ID/SCIM integration, rebuilt Collections, and MSP-level identity management.

└──▷ HOW TO FIND IT

MSPs managing multiple organizations can configure synced policy assignment for all orgs from one place instead of entering each org individually.

> 📍 From the MSP dashboard, navigate to Identities to create or manage Collections and Identity Connections across all organizations in a single view.

New Identities framework in the DNSFilter dashboard unifies users, collections, and policy assignment across identity sources for enterprise and MSP environments.

Direct Microsoft Entra ID integration via SCIM provisioning eliminates the AD Sync Tool dependency, with support for Entra Hybrid environments.

Rebuilt Collections engine supports large directories, multiple identity sources, and high-complexity group structures.

New unified Identities view gives administrators a consolidated picture of all identity sources, user objects, collections, and policy assignments.

MSP dashboard now supports managing Users, Collections, and Identity Connections across all organizations from a single view, without navigating into each organization individually.

WAS THIS USEFUL? 👍 👎

Zscaler Internet Access

Feedback

changelog · Jul 15, 2026 BLOG ↗

Zscaler Internet Access is a cloud-based secure web gateway that protects users by inspecting and filtering internet traffic to block threats and enforce policies.

Zscaler Endpoint Context enriches security logs with endpoint data for SIEM and SOC automation workflows.

WAS THIS USEFUL?

Zscaler Private Access

Feedback

changelog · Jul 15, 2026 BLOG ↗

Zscaler Private Access is a zero trust network access solution that securely connects users to private applications without exposing them to the network.

Zscaler adds Endpoint Context policy enforcement and Autonomous Application Shield for continuous private-app protection.

New Endpoint Context capability feeds endpoint-derived intelligence into Zero Trust Exchange policy decisions across TLS/SSL inspection, Zero Trust Firewall, DNS security, intrusion prevention, and Advanced Threat Protection.

New Autonomous Application Shield continuously assesses private application exposure via App Connectors and automatically applies app-specific protections in the Zscaler cloud to cover the window between vulnerability disclosure and patching.

WAS THIS USEFUL?

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ RegScale 6.31.1.5 (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 6.31.1.5 adds custom rules support for POA&M configuration.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

OpenAI Codex CLI

Feedback ⓘ rust-v0.145.0-alpha.16

3 RELEASES · 2026-07-15 → 2026-07-16 NOTES ↗ CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI alpha.16 adds prompt cache write token tracking and stronger forced `rm` command detection.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.145.0-alpha.16
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.16
```

Tracks prompt cache write token usage (`cache_write_input_tokens`) and reports it in analytics, metrics, traces, and rollout data alongside existing cache-read usage.

Strengthens forced `rm` command detection to catch forced removals inside control flow, substitutions, pipelines, traps, nested shells, and `sudo` / `env` wrappers — including combined, separate, and long-form force options even when they follow operands.

Extends image generation timeout in code mode to 120 seconds for both the initial call and subsequent waits.

Improves turn history summaries to use only `final_answer` `-phase` agent messages, excluding commentary items.

└──▷ BREAKING ON UPGRADE

! The `template_id` field is removed from MCP tool call items, lifecycle events, and `McpToolCallAppContext` app-server responses — any client or integration reading `template_id` (or `templateId`) from these payloads will no longer find it.

▾ 2 more releases in this issue · 2026-07-15 → 2026-07-16

rust-v0.145.0-alpha.15 NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI alpha.15 adds external agent memory migration, paginated subagent history, and smarter MCP capability-root propagation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.145.0-alpha.15
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.15
```

Adds external agent memory migration: discovers project memory Markdown files, preserves project scope, copies selected projects into the Codex memory extension workspace, and enqueues memory consolidation after workspace changes.

Propagates deferred environment capability roots to MCP, making executor plugins available as soon as their environment is ready and refreshing the MCP runtime when the selected root set changes.

Preserves paginated history for spawned or forked subagents, inheriting the parent's model context as a rollout prefix and resuming from stored model context instead of legacy rollout history.

Warns when conflicting capability root locations are registered under the same ID, retaining the first location and logging the conflict instead of silently dropping it.

Selects the elevated Windows sandbox automatically when network proxies are in use.

rust-v0.145.0-alpha.14 NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI adds Frameless Bidi realtime voice, MCP cache opt-out, Cursor import support, and paginated thread history.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.145.0-alpha.14
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.14
```

Adds realtime conversation protocol v3 with Frameless Bidi support, translating audio, transcripts, handoffs, session context, and lifecycle events over WebSocket and WebRTC.

Allows MCP servers to opt out of tool catalog caching via the experimental `codex/tool-catalog-cache` capability (`cacheable: false`), preventing stale snapshots from being served.

Adds Cursor IDE support to the setup import flow, joining existing editor integrations.

> Migrates plugin commands into skills automatically on install, converting legacy plugin command definitions to the Codex skill format.
>
> Scans skill roots concurrently (up to eight at a time) for faster startup on repos with many skills.
>
> +1 more

WAS THIS USEFUL? 👍 👎

Cline

Feedback ⓘ cli-v3.0.41 2 RELEASES · 2026-07-15 NOTES ↗

CODE ↗ LATEST STABLE

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

Cline CLI v3.0.41 adds compaction progress in the TUI, model ID suggestions from OpenAI-compatible endpoints, and persistent workspace git info.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch cli-v3.0.41 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout cli-v3.0.41
```

Shows compaction progress status in the TUI so you can track context-reduction operations in real time.

Suggests model IDs from OpenAI-compatible endpoints when configuring a provider, reducing manual lookup.

Persists and refreshes workspace git branch and remote info across sessions.

Blocks compaction from running during an active turn, preventing interference with in-progress responses.

▾ 1 more release in this issue · 2026-07-15

sdk/sdk/v0.0.61 NOTES ↗ STABLE

Cline SDK v0.0.61 adds context compaction progress, Plan/Act system-prompt guidance, and OpenAI-compatible model ID suggestions.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch sdk/sdk/v0.0.61
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout sdk/sdk/v0.0.61
```

Adds progress status reporting while context compaction runs, giving visibility into long-running compaction operations.

Persists and refreshes workspace git info (branch/remote) across sessions so agents maintain accurate repo context.

> Adds Plan/Act mode guidance to the system prompt, with nudges when switching between modes.
>
> Suggests model IDs from OpenAI-compatible endpoints when configuring custom providers.
>
> Centralizes provider-specific request headers in the LLM layer for consistent multi-provider header handling.
>
> +2 more

WAS THIS USEFUL? 👍 👎

OpenCode

Feedback ⓘ v1.18.2 NOTES ↗ CODE ↗

The open source coding agent.

Adds Mod+N shortcut for new tabs and a configurable `subagent_depth` limit to prevent runaway nested subagents.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.18.2 https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.18.2
```

Adds `Mod+N` as an alternative keyboard shortcut for opening a new tab in the desktop app.

Adds configurable `subagent_depth` limit to control how deeply subagents can spawn nested subagents (defaults to 1).

WAS THIS USEFUL? 👍 👎

Pi

Feedback ⓘ v0.80.8 NOTES ↗

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

Pi v0.80.8 adds unified model/auth runtime, live catalog refresh, and xAI Grok 4.5 with device-code OAuth.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.80.8 https://github.com/earendil-works/pi.git
# already have the repo? check out this version:
$ git checkout v0.80.8
```

└──▷ TRY IT

Force a full model catalog refresh after adding a new provider, without reinstalling pi or extensions.

```
$ pi update --models
```

Adds `ModelRuntime` as the canonical async SDK facade, centralizing model configuration, provider authentication, and dynamic provider catalogs.

Adds `pi update --models` command to force an immediate model catalog refresh without updating pi or extensions.

Adds `/model` live background refresh of configured providers, updating the selector with partial results as they arrive.

Adds provider-owned `/login` discovery, surfacing ambient auth status and informational links directly from registered providers.

Adds xAI device-code OAuth login so users can authenticate to xAI interactively without an API key.

+3 more

└──▷ BREAKING ON UPGRADE

! The `CreateAgentSessionOptions.authStorage` and `modelRegistry` SDK options are replaced by the async `modelRuntime` option; `AuthStorage` and its storage backends are no longer exported — use `ModelRuntime`, a custom pi-ai `CredentialStore`, or readStoredCredential() instead.

! ModelRuntime.getAll(), find(), getSnapshot(), and getAuthOptions() are removed; use pi-ai `Models` methods getModels(), getModel(), getProviders(), and checkAuth() directly.

! Request-auth assembly via ModelRegistry.getApiKeyAndHeaders() is replaced by ModelRuntime.getAuth(); passing a provider ID returns provider-scoped auth, passing a model also resolves built-in, `models.json`, and extension model headers.

! ModelRegistry.refresh() changed from synchronous `void` to `Promise<void>`; extensions must `await` it before making synchronous registry reads.

WAS THIS USEFUL? 👍 👎

Crush

Feedback ⓘ v0.85.0 NOTES ↗ CODE ↗

Glamourous agentic coding for all

Crush v0.85.0 adds an interactive question tool, an elapsed-time timer, and LM Studio vision support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.85.0
https://github.com/charmbracelet/crush.git
# already have the repo? check out this version:
$ git checkout v0.85.0
```

Adds a first-class question tool letting LLMs prompt users with single-choice, multiple-choice, free-form text, and yes/no forms — with full mouse support.

Displays an elapsed-seconds timer so you can see how long the LLM has been running.

Enables LM Studio vision capabilities (image support) via `SupportsImages`.

Adds automatic transparent auth retry so dropped credentials are refreshed without user intervention.

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ v2.1.211 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 v2.1.211 adds subagent stream visibility, cross-worktree permission persistence, and scientific-notation env vars.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.1.211 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.211
```

└──▷ TRY IT

Capture subagent reasoning in a CI pipeline that parses stream-json to audit what sub-tasks were run.

```
$ claude --output-format stream-json --forward-subagent-text -p 'Refactor the auth module and write tests' | tee agent-stream.ndjson
```

Override a token-budget environment variable using scientific notation — useful in shell configs where large numbers are error-prone.

```
$ CLAUDE_CODE_MAX_OUTPUT_TOKENS=1e6 claude -p 'Summarize the entire repo'
```

Adds `--forward-subagent-text` flag and `CLAUDE_CODE_FORWARD_SUBAGENT_TEXT` env var to surface subagent text and thinking inside `--output-format stream-json` output.

Changes `always allow` permission rules to save at the repository root, so approvals granted in one git worktree persist across all worktrees and sessions.

Accepts scientific notation and digit-separator spellings (e.g. `1e6`, `64_000`) for integer environment variables such as timeouts, token budgets, and retry counts.

Enables `save_to_disk` on Claude-in-Chrome screenshot actions to actually write the image to disk and return its path.

Improves background agent result reporting so Claude waits for real completion and reports still-running agents instead of fabricating results.

+3 more

WAS THIS USEFUL?  

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

PyPA python-pipx

Feedback ⓘ 1.15.1 NOTES ↗

Install and Run Python Applications in Isolated Environments

pipx 1.15.1 adds `outdated`, `exec`, `reset`, PEP 751 lock file support, scriptable JSON output, and major performance improvements.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 1.15.1 https://github.com/pypa/pipx.git
# already have the repo? check out this version:
$ git checkout 1.15.1
```

└──▷ TRY IT

Check which of your globally installed tools have newer versions available before deciding what to upgrade.

```
$ pipx outdated
```

Filter `pipx list` output to inspect a single installed package without scrolling through everything.

```
$ pipx list black
```

Start fresh after a corrupted pipx state — wipes all managed environments and reinitialises pipx.

```
$ pipx reset
```

New `pipx outdated` command lists all installed tools that have upgrades available.

New `pipx exec` command runs an already-installed application directly without re-resolving environments.

New `pipx reset` command wipes and rebuilds pipx's state from a clean install.

New `pipx cache` subcommand manages cached `pipx run` environments.

Adds PEP 751 lock file support for `pipx install`, enabling reproducible installs from lock files.

+17 more

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 16, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue