THE CYBER TOOLCHAIN

NO. 021

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 21, 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 33 tools

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Portkey ⓘ snapshot-20260721 NOTES ↗

Portkey is an AI gateway and observability platform that helps developers manage, monitor, and optimize LLM API calls with routing, caching, and reliability features.

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 API keys from your own vault at request time instead of storing them in Portkey — critical for orgs with strict secrets governance.

> 📍 In the Portkey console, go to your Virtual Key or Integration settings and add a `secret_mappings` block pointing to your AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault path. Portkey will resolve the credential at runtime.

Cap a team's weekly LLM spend and keep chat completions under a separate limit from embeddings calls.

> 📍 In the Portkey console, open Budget & Rate Limit Policies, set a `rpw` (requests-per-week) budget window for the relevant virtual key or config, then add an endpoint-scoped limit targeting the chat completions route independently from embeddings.

Route Bedrock guardrail evaluation to a private regional endpoint instead of the default public Bedrock URL.

> 📍 In the Portkey console, open the Bedrock guardrail plugin configuration and set `customHost` to your private or regional Bedrock-compatible endpoint URL.

- Introduces weekly rate and budget windows (`rpw`) so spend caps align with team planning cycles, not just 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.
- Enables GCP Workload Identity Federation log storage (`gcs_assume`) from AWS-hosted gateways via `GCP_WIF_AUDIENCE` and `GCP_WIF_SERVICE_ACCOUNT_EMAIL` , eliminating static GCP keys in AWS.
- Allows organization admins to include archived workspaces in analytics graphs and filters; saved filters now also accept workspace slugs alongside IDs.
- Integrates Zscaler AI Guard as a guardrails partner, applying Zscaler Detection Policies to LLM inputs and outputs via `beforeRequestHook` / `afterRequestHook` with a required `policyId` .

+6 more

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

DataDome ⓘ Envoy module 2.7.0 (2026-07-21) BLOG ↗

DataDome is a bot protection and fraud prevention platform that detects and blocks malicious traffic, scraping, and cyberattacks in real time.

DataDome Envoy module 2.7.0 exposes enriched headers as Envoy dynamic metadata for downstream logging.

› Adds DataDome enriched headers to Envoy dynamic metadata, making bot-detection signals available for use in access logs and observability pipelines.

WAS THIS USEFUL? 👍 👎

Cloudflare ⓘ WAF - WAF Release - 2026-07-21 (2026-07-21)

2 RELEASES · 2026-07-21 BLOG ↗ LATEST STABLE

Cloudflare is a cloud-based platform that provides website performance optimization, DDoS protection, CDN services, and application security to safeguard and accelerate internet traffic.

Cloudflare WAF adds new SSRF, LFI, XSS obfuscation, Adobe ColdFusion, and WordPress rules plus generic RCE/Auth Bypass/Info Disclosure detections.

- Adds two new SSRF detection rules blocking restricted-protocol and obfuscated-host request forgery attempts.
- Adds a new LFI path traversal detection rule to the Cloudflare Managed Ruleset.
- Adds three new XSS rules targeting JavaScript bracket-concatenation obfuscation across body, headers, and URI.
- Adds targeted blocking rules for Adobe ColdFusion file-upload path traversal (CVE-2026-48276) and directory-sequence path traversal (CVE-2026-48282).
- Adds blocking rules for WordPress unauthenticated SQL injection (CVE-2026-60137) in both the Cloudflare Managed and Free Rulesets.

+2 more

▾ **1 more release in this issue** · 2026-07-21

> WAF - WAF Release - Scheduled changes for 2026-07-27 (2026-07-21) BLOG ↗ STABLE
>
> Cloudflare WAF adds three new/improved managed rules for SSRF, command injection obfuscation, and generic DoS detection.
>
> - New detection rule for Command Injection - Obfuscation (rule ID ...761e7a4c), deploying in Log mode on 2026-07-27.
> - New detection rule for Generic Rules - DoS (rule ID ...90dcdb0a), deploying in Log mode on 2026-07-27.
> - Improved SSRF - Cloud - Beta detection rule (rule ID ...c2e84e2d), deploying in Log mode on 2026-07-27.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS ⓘ

Mention members and teams in incident notes (2026-07-20) BLOG ↗

GitGuardian SaaS is a secrets detection platform that scans code repositories and developer workflows to identify and remediate exposed credentials and sensitive data.

GitGuardian adds @mention notifications in incident notes, source health for new integrations, and oversized PR scanning on Business plan.

└──▷ HOW TO FIND IT

Loop in a teammate or an entire team on a critical incident without leaving GitGuardian.

> 📍 In the console, open any incident › add a note or feedback › type @ and select a member (@name) or team (@team-name) from the searchable list. The mentioned users receive an email linking directly to the incident.

Filter the incident list to only incidents from active sources, setting aside noise from archived or deleted ones.

> 📍 In the console, go to the Perimeter page or Incident list › open the source status filter › select Monitored (or exclude Archived / Deleted) › save as a view for reuse.

Reduce false positives on a Custom Source install by limiting scans to detectors that can verify secret validity.

> 📍 In the console, go to the Custom Source install settings › enable 'use only detectors with validators' › save. Subsequent scans on that source will drop matches whose detector has no validator before the validity check runs.

- Adds @mention support in incident notes and feedback — type @ to tag a member by @name or an entire team by @team-name, triggering email notifications that link directly to the incident.
- Extends source health management (auto-pause, auto-resume, and recovery guidance) to ServiceNow, JFrog Container Registry, JFrog Package Registry, and Microsoft OneDrive sources.
- Introduces per-source status indicators (Monitored, Unmonitored, Unreachable, Archived, Deleted) on the perimeter page and incident list, with filterability and saveable views.
- Enables scanning of oversized pull requests (>200 commits, >60,000 files, or >300,000 lines changed) via repository cloning on the Business plan, bypassing GitHub API rate

limits.

› Adds opt-in 'use only detectors with validators' option for Custom Source (BYOS) installs, dropping unverifiable matches before validity checks to reduce noise.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

EMBA ⓘ v2.0.3-legacy-time NOTES ↗

EMBA - The firmware security analyzer

EMBA v2.0.3 adds legacy SSL/TLS support to reach end-of-life devices modern OpenSSL can no longer connect to.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.0.3-legacy-time https://github.com/e-m-b-a/emba.git
# already have the repo? check out this version:
$ git checkout v2.0.3-legacy-time
```

› Supports legacy SSL/TLS ciphers during live analysis, enabling vulnerability detection on old embedded devices that modern OpenSSL rejects.

› Expands CSV log output across additional modules for richer machine-readable reporting.

› Adds entropy file checking to aid detection of encrypted or compressed firmware regions.

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

malcontent ⓘ v1.25.4 NOTES ↗

#supply #chain #attack #detection

malcontent v1.25.4 adds `--ignore-rules` to suppress specific YARA rules from scan and diff reports.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.25.4 https://github.com/chainguard-dev/malcontent.git
# already have the repo? check out this version:
$ git checkout v1.25.4
```

└──▷ TRY IT

Suppress noisy or known-false-positive rules by name during a scan so they don't clutter results.

```
$ mal scan --ignore-rules my_noisy_rule,another_rule <target>
```

Exclude a specific rule when diffing two container images to focus on net-new risky behaviors.

```
$ mal diff --ignore-rules some_rule --image cgr.dev/chainguard/myapp:v1.0.0 cgr.dev/chainguard/myapp:v1.0.1
```

› Adds `--ignore-rules` flag to exclude specific YARA rules from scan and diff report output.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

mac_apt ⓘ v1.33.2 NOTES ↗

macOS (& ios) Artifact Parsing Tool

mac_apt v1.33.2 adds Biome and Trash plugins, compressed DMG support, SFL4 in RECENTITEMS, and richer bookmark metadata.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.33.2 https://github.com/ydkhatri/mac_apt.git
# already have the repo? check out this version:
$ git checkout v1.33.2
```

- New BIOME plugin parses Apple Biome data from macOS artifacts.
- New TRASH plugin extracts deleted file/folder metadata from .`DS_S`tore files in the .Trash folder.
- Supports compressed DMG images (based on dmgwiz), expanding disk image compatibility.
- Adds SFL4 processing to the RECENTITEMS plugin for broader recent-items coverage.
- Extracts volume info and item creation date from RECENTITEMS bookmarks.

WAS THIS USEFUL? 👍 👎

dfvfs ⓘ 20260717 NOTES ↗

Digital Forensics Virtual File System (dfVFS)

dfvfs 20260717 switches to pyfsfat as the default FAT file-system back-end.

└──▷ GET THIS VERSION

```
$ git clone --branch 20260717
https://github.com/log2timeline/dfvfs.git
# already have the repo? check out this version:
$ git checkout 20260717
```

› Promotes pyfsfat to the default back-end for FAT file-system parsing, replacing the previous default.

└──▷ BREAKING ON UPGRADE

! pyfsfat is now the default FAT back-end; setups relying on the previous default FAT back-end behavior may see differences in FAT volume parsing on upgrade.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

NVIDIA OpenShell ⓘ v0.0.88 NOTES ↗

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

OpenShell v0.0.88 adds a workspace resource model with scoping and membership controls.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.0.88 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.88
```

› Adds a workspace resource model with scoping and membership support, enabling structured multi-tenant workspace management.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Anchore syft ⓘ v1.49.0 NOTES ↗

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

Syft v1.49.0 adds support for OCI image index manifests in root OCI layouts.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.49.0 https://github.com/anchore/syft.git
# already have the repo? check out this version:
$ git checkout v1.49.0
```

› Supports `application/vnd.oci.image.index.v1+json` manifests in root OCI layout directories, enabling SBOM generation from multi-platform OCI index artifacts.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Trend Vision One ⓘ

Extend file password analysis to all inbound security functions (2026-07-21)

4 RELEASES · 2026-07-16 → 2026-07-21 BLOG ↗ LATEST STABLE

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

File Password Analysis now covers all inbound security functions—spam, content filtering, and DLP—not just virus scanning.

└──▷ HOW TO FIND IT

Enable File Password Analysis under your inbound policy so decrypted attachment content is automatically scanned by spam filtering, content filtering, and DLP—closing the gap attackers exploit with encrypted attachments.

📍 In the console, go to Email and Collaboration Security → Cloud Email Gateway Protection → Inbound Protection, then enable File Password Analysis in your inbound policy.

› Extends File Password Analysis to all inbound security functions, including spam filtering, content filtering, and DLP, so password-protected attachments are fully inspected across every enabled policy layer.

▾ 3 more releases in this issue · 2026-07-16 → 2026-07-21

Agentless Vulnerability & Threat Detection version 1.636.06121414 - AWS eu-south-2 (Spain) Region support (2026-07-20)

BLOG ↗ STABLE

Agentless Vulnerability & Threat Detection now supports AWS eu-south-2 (Spain) region for scanning and VPC Flow Log detections.

└──▷ HOW TO FIND IT

Enable coverage for the new Spain region by updating your existing CloudFormation stack via Cloud Accounts settings — no redeployment required.

📍 In the console, go to Cloud Security → Cloud Accounts → AWS and perform a stack update on your existing deployment to enable the eu-south-2 (Spain) region.

› Supports deploying Agentless Vulnerability & Threat Detection in the AWS eu-south-2 (Spain) region.

› Supports Cloud Detections for AWS VPC Flow Logs in the AWS eu-south-2 (Spain) region.

On-premises gateway supports passing user information to upstream proxy (2026-07-20)

BLOG ↗ STABLE

Zero Trust Secure Access on-premises gateway can now stamp user identity into the X-Authenticated-User header for upstream proxies.

└──▷ HOW TO FIND IT

Propagate authenticated user identity to an upstream proxy so it can enforce per-user policies or logging.

📍 In the console, go to Zero Trust Secure Access → Secure Access Configuration → Internet Access and AI Secure Access Configuration and enable the option to pass user information to the upstream proxy.

› Enables the on-premises gateway to pass user identity to an upstream proxy via the X-Authenticated-User header.

Agentic red-teaming with AI Scanner (2026-07-16)

BLOG ↗ STABLE

Trend Vision One AI Scanner adds LLM-driven agentic red-teaming that dynamically adapts attack prompts based on target AI app responses.

└──▷ HOW TO FIND IT

Configure and run an agentic attack scan to continuously validate AI application defenses against adaptive, LLM-driven attack sequences.

📍 In the console, go to AI Security → AI Application Security → TrendAI Artifact Scanner, configure an agentic attack by defining your attack objective, then run it alongside your existing attack presets or custom prompts.

› New agentic attacks capability in AI Scanner uses an LLM-driven attacker model that generates and iterates attack prompts dynamically based on target AI application responses, rather than a static prompt library.

› Agentic attacks are configured through the TrendAI Artifact Scanner and can run alongside existing attack presets, attack objectives, and custom prompts in the same scan.

› Enables continuous red-team validation of AI defenses — guardrails, system prompts, and model behavior — against evolving, adaptive adversarial techniques.

WAS THIS USEFUL? 👍 👎

Automox ⓘ

Now live: Partial Policy Success and Per-Package Visibility in the Policy Results Report (2026-07-20)

BLOG ↗

Automox is a cloud-native endpoint management platform that automates patching, configuration, and software deployment across Windows, macOS, and Linux devices.

Automox Policy Results Report gains Partial status, per-package failure reasons, and raw log side panel.

.
.

└──▷ HOW TO FIND IT

Spot runs that need attention after a patch cycle by filtering the Policy Results Report for 'Partial' outcomes, then expanding a device row to see which package failed and why.

📍 In the Policy Results Report, use the Results filter and select 'Partial' to isolate mixed-outcome runs. Expand any device row to view per-package results and failure reasons. Click 'View Raw Logs' on a failed package to open the full log in the side panel.

- Adds 'Partial' as a new policy run result status when some patches succeed and others encounter errors, replacing a single pass/fail outcome.
- Adds per-package breakdowns in the device table for Successful, Failed, or Partial runs, showing plain-language failure reasons (e.g., 'The application is currently running', 'Pending system restart').
- Raw logs for individual packages now open in a side panel via 'View Raw Logs', keeping detail accessible without cluttering the table.
- Replaces the stacked bar chart and legend in the Results column with a direct count of each status, with definitions moved to filter tooltips.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines ⓘ Track cases with change control in the storyboard (2026-07-20)

3 RELEASES · 2026-07-20 BLOG ↗ LATEST STABLE

Tines is a no-code automation platform that helps security and operations teams build workflows to automate repetitive tasks and incident response processes.

Tines extends change control to cover all case elements defined in the storyboard.

› Enables change control tracking for every change across case elements defined in the storyboard.

▾ **2 more releases in this issue** · 2026-07-20

Formula support for case details (2026-07-20) BLOG ↗ STABLE

Tines cases now support formulas across every case property, primitive, and case content field from the storyboard.

› Enables formula expressions across all case properties, primitives, and case content when defining cases from the storyboard.

Visually create cases in the storyboard (2026-07-20) BLOG ↗ STABLE

Tines lets you visually configure case creation and write rich-text descriptions directly in the storyboard editor.

└──▷ HOW TO FIND IT

When building a case-creation workflow, configure all case fields and format the description visually without leaving the storyboard.

> 📍 In the storyboard, add a Create Case tile, then use the visual editor to define case fields and apply rich-text formatting to the description and block content.

› Enables visual definition of all case fields when building with the create case tile in the storyboard.

› Supports rich-text formatting for case descriptions and block content within the storyboard editor.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Qualys Patch Management ⓘ changelog · Jul 21, 2026 BLOG ↗

Qualys Patch Management is a cloud-based tool that automates the discovery, prioritization, and deployment of patches across IT assets to reduce vulnerabilities.

Qualys TruRisk Eliminate adds AI-Powered Patch Reliability Score and zero-touch automation for low-risk application patching.

.
.
.
.

└──▷ HOW TO FIND IT

Use Patch Now to immediately deploy patches for a specific CVE directly from the vulnerability finding in Qualys Patch Management.

> 📍 In the Qualys Patch Management console, locate CVE-2026-8933 in the vulnerability list, then click 'Patch Now' (to the right of the vulnerability) to add it to a patch job targeting vulnerable assets.

- New AI-Powered Patch Reliability Score uses crowd-sourced signals from patch deployments across the internet to assess whether a patch is safe to deploy before rollout.
- New zero-touch patch automation enables low-risk applications (e.g., browsers and standalone software) to receive patches automatically after a one-time job setup, with no manual intervention required.
- Supports autonomous remediation workflow: high-reliability patches deploy with confidence while low-reliability patches are flagged for review, balancing speed and production safety.

WAS THIS USEFUL? 👍 👎

Qualys VMDR ⓘ changelog · Jul 21, 2026 BLOG ↗

Qualys VMDR is a cloud-based vulnerability management platform that continuously discovers, assesses, prioritizes, and remediates security vulnerabilities across IT assets.

Qualys VMDR adds QID coverage to detect CVE-2026-8933, a high-severity LPE in snap-confine on Ubuntu 24.04/25.10/26.04.

└──▷ HOW TO FIND IT

Search the Qualys Vulnerability Knowledgebase for the new QIDs to verify coverage and prioritize patching of affected Ubuntu systems.

> 📍 In the Qualys VMDR console, go to Vulnerabilities › Knowledgebase and search for CVE-2026-8933 to retrieve the associated QIDs and confirmed detection status.

› New QID detection(s) for CVE-2026-8933, a high-severity Local Privilege Escalation in snap-confine affecting default Ubuntu Desktop 24.04, 25.10, and 26.04 installations.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Intruder ⓘ

Microsoft 365 security checks for Azure integrations (2026-07-21) BLOG ↗

Intruder is a cloud-based vulnerability scanner that continuously monitors and identifies security weaknesses in internet-exposed systems, APIs, and web applications.

Intruder now surfaces Microsoft 365 configuration risks across Entra ID, Exchange Online, SharePoint, and Teams within Azure cloud scans.

└──▷ HOW TO FIND IT

Unlock full M365 security coverage by updating your existing Azure integration's permissions in the Intruder console.

> 📍 In the console, go to your Azure integration settings and update the integration's permissions following the steps outlined in the M365 setup guide to enable the full set of Microsoft 365 checks.

- Adds Microsoft 365 configuration security checks covering Entra ID, Exchange Online, SharePoint, and Teams as part of existing Azure cloud security scans.
- Automatically runs a subset of M365 checks with existing Azure integrations, with full coverage unlocked by updating integration permissions.

WAS THIS USEFUL? 👍 👎

Censys

July 20, 2026 (2026-07-20) BLOG ↗

Censys is an internet intelligence platform that continuously scans and indexes internet-connected devices, helping security teams discover and monitor exposed assets and vulnerabilities.

Censys adds Argo CD scanner, three AI tool fingerprints, and three Microsoft SharePoint RCE risks to ASM.

└──▷ TRY IT

Identify Microsoft SharePoint servers at risk of unauthenticated RCE (CVE-2026-50522) in your ASM workspace.

```
$ risks.name: "Vulnerable Microsoft SharePoint Server [CVE-2026-50522]"
```

Surface SharePoint servers flagged for the CISA KEV-listed authenticated deserialization RCE (CVE-2026-45659) across your inventory.

```
$ risks.name: "Microsoft SharePoint Server Authenticated Deserialization RCE [CVE-2026-45659]"
```

- Adds new endpoint scanner for Argo CD (ARGOCD), available to Starter, Search, and Core users.
- Adds fingerprints for three AI/no-code tools: Anything LLM, Capacity, and Softr.
- Adds three new ASM risks for Microsoft SharePoint Server covering CVE-2026-45659 (authenticated deserialization RCE, CISA KEV-listed), CVE-2026-50522, and CVE-2026-58644 (both unauthenticated deserialization RCE).
- Publishes Censys ARC Rapid Response advisory for Microsoft SharePoint Server RCE [CVE-2026-50522, CVE-2026-58644] with platform, ASM, and legacy search queries to identify exposed instances.

WAS THIS USEFUL? 👍 👎

Axonius

Axonius Release Notes 9.0.1 (2026-06-25)

BLOG ↗

Axonius is a cybersecurity asset management platform that discovers, inventories, and manages all devices, users, and cloud assets across an organization's IT environment.

Axonius 9.0.1 adds 5 new adapters, SAML advanced role/scope mapping, custom compliance values, and dozens of adapter enhancements.

└──▷ HOW TO FIND IT

Assign roles from one SAML attribute (e.g., group membership) and data scopes from another (e.g., department) for fine-grained SSO access control.

> 📍 In the console, go to Settings › Access Settings › SAML, enable Advanced Mapping, then configure the Role mapping table with your group-membership claim and the Data Scope mapping table with your department claim.

Override Axonius default compliance benchmarks with your organization's own thresholds for SaaS application settings.

> 📍 In the console, go to SaaS Applications › Settings, locate the setting to customize, and set a Custom Setting Recommended Value to define your compliance standard.

Give all dashboard viewers a consistent default filter view without overriding individual saved preferences.

> 📍 In the console, open a Dashboard, click Saved Filters, hover over the target filter, click the 3-dot menu, and select Set as Default.

- New Bold Security adapter: inventories endpoints, agent status, and browser extension protection.
- New EasyDNS adapter: fetches domain registrations, zone management, and DNS record configurations.
- New Google Threat Intelligence ASM adapter: discovers internet-facing assets including domains, networks, network services, and aggregated security findings.
- New Lenovo Device Orchestration adapter: fetches devices and users with AI-driven fleet health, hardware diagnostics, and warranty data.
- New One Identity Safeguard adapter: fetches managed devices with privileged access management context including credential vaulting and session monitoring.

+34 more

└──▷ BREAKING ON UPGRADE

- ! HPE Aruba Networking ClearPass Policy Manager: the Extended Info Fetch Window (Days) value is now constrained to 0–90; existing configurations with values outside this range must be updated.
- ! Microsoft Defender for Endpoint: the default value for vulnerability aggregation has changed to True, which may alter existing fetch behavior on upgrade.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter ⓘ

07/21 July 21 – Web Application and API Update (2026-07-21) BLOG ↗

DNSFilter is a cloud-based DNS security platform that blocks malicious websites, filters unwanted content, and protects networks from cyber threats in real time.

DNSFilter redesigns the Query Log with multi-select filters, cross-org toggling, second-level timelines, and shareable URLs.

└──▷ HOW TO FIND IT

Share a precise cross-org Query Log investigation with a colleague by bookmarking the URL with active filters intact.

> 📍 In the Query Log, select the target Sites, Roaming Clients, and Users via the filter bar, set a custom timeline down to the second, then copy the URL from your browser — paste it to a colleague or save it as a bookmark to restore the exact view.

- Redesigned DNS Query Log with a top-level filter bar supporting multi-select across Sites, Roaming Clients, Relays, and Users.
- MSP admins can now toggle between organizations directly from the Query Log without losing active filters, enabling faster cross-org investigation.
- New custom timeline picker lets analysts narrow Query Log search windows down to the second.
- Active filters persist in the URL, making Query Log views bookmarkable and shareable.
- New Query Log data columns added: Resolved IPs, Query Type, and Roaming Client OS Type.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Ping Identity ⓘ 20 Jul 2026 (2026-07-21) BLOG ↗

Ping Identity is an identity security platform that provides authentication, single sign-on, multi-factor authentication, and access management solutions for enterprises.

Ping Identity adds Rich Authorization Requests (RFC 9396) support, per-client authorization detail types, and a new Patch Object node v2.0.

└──▷ HOW TO FIND IT

Restrict which authorization detail types a specific OAuth 2.0 client may request, scoping RAR to only the types relevant to that client.

> 📍 In the Admin console, go to Applications › OAuth 2.0 Clients › <your client> › Advanced and set the 'Authorization Detail Type(s)' field to the permitted types, then save.

- New realm-level OAuth2 Provider setting 'Enable Rich Authorization Requests with RCS' enables processing of `authorization_details` requests when an RCS is configured and controls advertisement of `authorization_details_types_supported` on the well-known endpoint.
- Supports configuring Authorization Detail Type(s) on individual OAuth 2.0 clients to control which RFC 9396 Rich Authorization Request detail types each client can request.
- Dynamic client registration requests can now include `authorization_details_types` to declare the authorization detail types a client supports.
- Adds the ability to skip signature verification when importing SAML metadata.
- New v2.0 of the Patch Object node always updates the managed identity object as Advanced Identity Cloud, removing the 'Patch as Object' configuration option (v1.0 is deprecated but remains functional in existing journeys).

WAS THIS USEFUL? 👍 👎

SailPoint ⓘ

Enhancement: Support for Stronger Integrations (2026-07-20)

2 RELEASES · 2026-07-20 BLOG ↗ LATEST STABLE

SailPoint is an identity security platform that helps organizations manage and govern user access to applications, data, and systems across the enterprise.

SailPoint ISC admins must now specify an approved integration when creating API or Personal Access Tokens.

- Requires admins to select an approved integration when creating API or Personal Access Tokens, enabling per-integration traffic visibility.
- Introduces the Unified Platform Access Program, enabling deeper product integrations tied to specific tokens.

└──▷ BREAKING ON UPGRADE

! When creating an API or Personal Access Token, admins must now specify the approved integration — existing token-creation workflows that do not supply an integration will require updating.

▾ **1 more release in this issue** · 2026-07-20

New Capability: Personal Access Token Management for Admins (2026-07-20)

BLOG ↗ STABLE

SailPoint admins can now view, edit, and bulk-manage all Personal Access Tokens org-wide, plus a new Bulk update PAT expiration API.

└──▷ HOW TO FIND IT

Audit and revoke stale PATs across your org before they become a security risk — reach this without contacting individual users.

📍 In the Admin console, go to Admin → Security Settings → Personal Access Tokens to view all tokens in the environment, then select one or more tokens and edit their expiration dates in bulk.

- New admin screen under Admin → Security Settings → Personal Access Tokens lets admins view and manage all PATs in the environment, not just their own.
- Admins can edit individual PATs or modify multiple tokens at once using bulk operations.
- Admins can change or remove expiration dates on any PAT across the organization.

› New Bulk update PAT expiration API enables programmatic management of PAT expiration dates at scale.

└──▷ BREAKING ON UPGRADE

! Any Personal Access Token created before February 5, 2026 that has never had an expiration date added or removed will automatically be assigned an expiration date 6 months (180 days) from the release date (production rollout: 7/28/2026), which may break long-lived integrations relying on those tokens remaining non-expiring.

WAS THIS USEFUL? 👍 👎

JumpCloud ⓘ

JumpCloud Agent Release Notes for July 20, 2026 (2026-07-20) BLOG ↗

JumpCloud is a cloud-based directory platform that manages user identities, devices, and access to applications and networks across organizations.

JumpCloud Agent 2.147.7 adds SMS OTP MFA support across all OSes and reduces SIEM noise on Linux.

› Adds SMS OTP authentication factor support for MFA on logins and password changes across all managed device OSes.

› Self Service App Catalog now processes state changes only for applications the user interacts with, eliminating unnecessary retries and speeding up app management.

› Removes unnecessary sudo usage from Linux agent operations, reducing SIEM security alerts and improving system audit log clarity.

WAS THIS USEFUL? 👍 👎

Auth0 ⓘ Anonymous sessions are now in beta! (2026-07-20) BLOG ↗

Auth0 is a cloud-based identity platform that provides authentication and authorization services, enabling developers to securely add login and access control to applications.

Auth0 launches Anonymous Sessions beta: stateless guest sessions with seamless linking to authenticated users

└──▷ HOW TO FIND IT

Enable anonymous sessions for a specific application and API, and set session length, via the Auth0 dashboard.

> 📍 In the Auth0 dashboard, go to Applications › <your application> and enable Anonymous Sessions. Then go to APIs › <your API> and enable Anonymous Sessions. Finally, set the session length under Tenant Settings › Anonymous Sessions.

- New `POST /anonymous/token` endpoint creates, refreshes, and retrieves access tokens for unauthenticated (anonymous) sessions.
- Supports seamless session linking — anonymous sessions can be tied to authenticated users via pre-registration and post-login Actions.
- Activity Preservation carries anonymous session data across devices onto the authenticated user profile upon login.
- New tenant-level setting to configure anonymous session length.
- Per-application and per-API configuration options to enable anonymous sessions.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Pi ⓘ v0.81.0 NOTES ↗

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

Pi v0.81.0 adds local llama.cpp model management, full provider extensions, Qwen Token Plan support, and expanded usage accounting.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.81.0 https://github.com/earendil-works/pi.git
# already have the repo? check out this version:
$ git checkout v0.81.0
```

└──▷ TRY IT

Connect to a local llama.cpp router and search for a Hugging Face model to download and load — without leaving the Pi session.

```
$ /login
/llama <model-search-query>
```

- Adds built-in llama.cpp router support via `/login` connection setup and `/llama` subcommand for Hugging Face model search, download, explicit load/unload, and live progress.
- Enables extensions to register complete pi-ai providers with native authentication, model refresh, filtering, and custom streaming behavior.
- Adds Qwen Token Plan and Qwen Token Plan China as built-in providers with regional endpoints and API-key authentication.
- Expands usage accounting to persist tool, compaction, and branch-summary usage in session totals and statistics.
- Adds `get_available_thinking_levels` RPC command and RpcClient.getAvailableThinkingLevels() method for querying model thinking-level options.

+1 more

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ rust-v0.145.0-alpha.29 3 RELEASES · 2026-07-21

NOTES ↗ CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Adds route-aware HTTP client pooling and system proxy support across daemon, plugins, and remote backends.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.145.0-alpha.29
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.29
```

- Adds a route-aware HTTP client pool (`RouteAwareClientPool`) that selects the correct outbound proxy route per request URL, reusing clients by resolved route with a bounded cache of 16 routes.
- Honors system proxy settings in the daemon updater, remote plugins, and cloud environment discovery.
- Supports `CARGO_HTTP_CAINFO` as a custom CA source in managed proxy/MITM environments.
- Makes external session detection limits configurable via `maxSessionAgeDays` and `maxSessions` in the external-agent config detection request.
- Adds optional `providerId` to `externalAgentConfig/import` for analytics attribution of external agent imports by provider.

+3 more

▾ 2 more releases in this issue · 2026-07-21

> rust-v0.145.0-alpha.28 NOTES ↗ CODE ↗ PRE-RELEASE
>
> Codex CLI gains concurrent terminal stdin, audio input/output, Amazon Bedrock custom transports, and incremental Markdown streaming.
>
> └──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.145.0-alpha.28
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.28
```

- Enables parallel `write_stdin` tool calls across independent terminal sessions, serializing per-session I/O for safety.
- Supports custom transports for Amazon Bedrock provider connections.
- Adds audio output support to dynamic tools and code mode.
- Adds occurrence search for paginated thread history.
- Adds batched executor capability discovery.

+13 more

rust-v0.145.0-alpha.27 NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI adds Esc key to interrupt running sessions and updates bundled model metadata including GPT-5.6 context windows.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.145.0-alpha.27
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.27
```

- Enables pressing Esc to interrupt a running agent session mid-task, with Esc now documented as an interrupt key in the status indicator.
- Updates bundled GPT-5.6 model variants with 272,000-token context windows and refreshed model instructions.
- Adds personality instruction variables for GPT-5.6 model variants and includes a GPT-5.5 availability notice in the bundled catalog.

└──▷ BREAKING ON UPGRADE

! The `auto_review` and `permissions` message fields and the legacy `supports_reasoning_summary_parameter` flag are removed from the bundled model catalog.

WAS THIS USEFUL? 👍 👎

OpenHands ⓘ cloud-1.47.0 NOTES ↗

OpenHands: AI-Driven Development

OpenHands cloud-1.47.0 gates Agent Canvas behind SaaS authentication.

└──▷ GET THIS VERSION

```
$ git clone --branch cloud-1.47.0
https://github.com/OpenHands/OpenHands.git
# already have the repo? check out this version:
$ git checkout cloud-1.47.0
```

› Protects the Agent Canvas behind SaaS authentication, restricting access to authenticated cloud users.

WAS THIS USEFUL? 👍 👎

Claude Code ⓘ v2.1.216 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 v2.1.216 adds sandbox filesystem-isolation bypass, smarter `/context` warnings, and improved `/fork`, `/ultrareview`, and `/code-review` ultra feedback.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.1.216 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.216
```

└──▷ USE IT

Allow a sandboxed session to read/write the filesystem freely while still blocking unexpected network egress — useful when your task requires broad file access but you still want outbound network control.

```json
# In your Claude Code settings (e.g. .claude/settings.json)
{
  "sandbox": {
    "filesystem": {
      "disabled": true
    }
  }
}
```

- Adds `sandbox.filesystem.disabled` setting to skip filesystem isolation while keeping network egress control active.
- Adds explicit context-window overflow warning to `/context` and surfaces failed `/compact` as a visible error.
- Improves `/fork` confirmation to include the new session name, `claude attach` id, and whether the copy shares the checkout.
- Improves `/ultrareview` diff-too-large error to display configured limits, measured diff size, and the largest contributing files.
- Improves `/code-review ultra` empty-diff message to name the exact base ref and suggest passing an explicit base.

+3 more

WAS THIS USEFUL? 👍 👎

Crush ⓘ v0.86.0 NOTES ↗ CODE ↗

Glamourous agentic coding for all

Crush v0.86.0 adds LSP rename/replace tools, scrollable sidebar, removable attachments, and SSH-friendly OAuth URL copy.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.86.0
https://github.com/charmbracelet/crush.git
# already have the repo? check out this version:
$ git checkout v0.86.0
```

.
.
.
.

└──▷ TRY IT

Get aggregate token/cost stats across all sessions or a specific directory tree without opening the TUI.

```
$ crush stats --all
```

Scope stats to a particular project directory when working in a monorepo.

```
$ crush stats --crawl-dir ./services/auth
```

- Adds `lsp_rename` tool: renames a symbol and all its references across the project via LSP.
- Adds `lsp_replace_symbol` tool: replaces, inserts before/after, or deletes an entire function, method, or class by name via LSP.
- Enables scrollable sidebar with keyboard (l/h/j/k/g/G/arrows) and mouse-wheel navigation when sidebar content overflows.
- Adds clickable ✕ button on attachment chips to remove attachments with the mouse.
- Adds <kbd>u< `/kbd` > keybinding in the OAuth dialog to copy the verification URL — enables authentication over SSH where auto-browser-open is unavailable.

+2 more

WAS THIS USEFUL? 👍 👎

OpenCode ⓘ v1.18.4 NOTES ↗ CODE ↗

The open source coding agent.

OpenCode v1.18.4 adds adaptive thinking for Kimi models and desktop UI improvements including inline file browser tabs.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.18.4
https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.18.4
```

- Adds adaptive thinking controls for Kimi models on Anthropic-compatible providers, with summarized reasoning output by default.
- Syncs the embedded terminal theme with the active app theme in the desktop app.
- Rewrites the v2 prompt input for more reliable command, context, shell, attachment, and history interactions.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

Hermes ⓘ v2026.7.20 NOTES ↗ CODE ↗

The agent that grows with you

Hermes v2026.7.20 slashes first-token latency 80%, adds password-manager secret sources, live subagent transcripts, and terminal subscription management.

└──▷ GET THIS VERSION

```
$ git clone --branch v2026.7.20
https://github.com/NousResearch/hermes-agent.git
# already have the repo? check out this version:
$ git checkout v2026.7.20
```

└──▷ TRY IT

Export a session history as a Hugging Face-ready trace with secrets scrubbed — useful for building fine-tuning datasets from real agent conversations.

```
$ hermes sessions export --redact
```

Manage your Nous subscription tier without leaving the terminal — view your plan, preview upgrade cost, and apply changes.

```
$ /subscription
```

Watch a subagent work in real time by tailing its live transcript as soon as it launches.

```
$ tail -f <transcript-file>
```

- Cuts cold-start first-turn latency ~80% (4.3s → 0.9s) across CLI, TUI, gateway, desktop, and cron.
- Streams reasoning model thinking live by default (`display.show_reasoning` on) — no more waiting at a spinner.
- Adds `/subscription` and `/topup` commands to manage your Nous plan and billing directly from the terminal or TUI.
- Adds `/deny <reason>` command to tell the agent why a command was refused so it can course-correct.

› Introduces user-defined deny rules that block commands even under yolo mode.

+15 more

WAS THIS USEFUL? 👍 👎

CrewAI ⓘ 1.15.5 NOTES ↗

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

CrewAI 1.15.5 adds authentication support for skill registry downloads.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.15.5 https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.5
```

› Supports authenticated downloads from the skill registry, enabling access to private or restricted skills.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ToolHive ⓘ v0.40.0 NOTES ↗

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

ToolHive v0.40.0 adds RFC 8693 token exchange for agentic delegation and early MCP 2026-07-28 revision support.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.40.0 https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.40.0
```

- Adds RFC 8693 OAuth 2.0 Token Exchange grant handler to the embedded authorization server, enabling agents to exchange a user token for a delegated token carrying both the user (`sub`) and the acting agent (`act.sub`) — the foundation for agentic delegation.
- Adds request-parsing support for the upcoming MCP 2026-07-28 ('Modern') stateless revision, including a revision classifier, `Mcp-Method` / `Mcp-Name` headers, `_meta` vocabulary, and `server/discover` / `subscriptions/listen` authz registration (dormant until wired into proxy routing).
- Migrates MCP protocol handling from `mark3labs/mcp-go` to the official `modelcontextprotocol/go-sdk` via the `toolhive-core/mcpcompat` compatibility shim.

└──▷ BREAKING ON UPGRADE

! The stdio bridge now forwards only progress/message notifications; `tools/list_changed` and similar notifications are dropped, so dynamic-capability servers may show stale capability lists until clients re-list.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Syslifters SysReptor ⓘ 2026.58 NOTES ↗

A customizable and powerful penetration testing reporting platform for offensive security professionals. Simplify, customize, and automate your pentest reports with ease.

SysReptor 2026.58 adds plugin management UI, SSO identity API token auth, and AI agent delete confirmation.

└──▷ GET THIS VERSION

```
$ git clone --branch 2026.58
https://github.com/Syslifters/sysreptor.git
# already have the repo? check out this version:
$ git checkout 2026.58
```

- Enables enabling/disabling plugins directly from the settings web UI.
- Supports SSO identity API for API token authentication.
- AI agent now prompts for confirmation before reverting changes (delete).

└──▷ BREAKING ON UPGRADE

! Forgot password functionality via email is disabled when `ALLOWED_HOSTS` is unset.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe