THE CYBER TOOLCHAIN

NO. 060

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 31, 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

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 159 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 39 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

Three unrelated things stand out. GitGuardian SaaS now scans S3 buckets, including inside compressed archives, so leaked credentials in

vim ? help

dumps and backups are findable without Git. NodeZero added an assumed-breach Azure pentest driven from a managed identity, testing what a compromised workload actually reaches. And Chainguard's Guardener enforces SHA-pinned Actions and signed commits directly in the repo.

DETECT

Hunt for secrets outside your Git history

Secret scanners have historically stopped at the repo boundary, which is exactly where dumps, logs and backups pile up; GitGuardian now walks S3 buckets and unpacks compressed archives, and AI-agent verdicts replace the old ML risk score for judging public leaks. Varonis adds AWS Organization-scoped targeting and scan-pause windows so a full-account crawl can be aimed at the right accounts and kept out of business hours.

GitGuardian SaaS · Varonis

IDENTIFY

See what an attacker inside your Azure tenant can actually reach

Cloud pentesting usually starts from a stolen key on somebody's laptop; NodeZero's new test type starts from a managed identity, which is the credential a compromised workload really holds, and a Cloud Connections flow handles authorizing the tenant and subscription. Paired with the Entra ID AD Password Audit, you get evidence of reachable paths rather than a config-drift list.

Horizon3 NodeZero

PROTECT

Make supply-chain hygiene a merge-time gate instead of a policy doc

Guardener installs as a GitHub App and does the tedious parts nobody gets to: rewriting Dockerfiles to hardened bases, pinning Actions to SHAs so a moved tag cannot swap code under you, and requiring signed commits, plus open-beta Rego pull policies for custom rules. JFrog's Platform Federation keeps curation policies and access control in sync peer-to-peer across self-managed and SaaS sites, so a blocked package stays blocked at every site without a per-site change.

Chainguard · JFrog

RESPOND

Resume a long scan or investigation instead of restarting it

> dirsearch's JSON session save/resume means a multi-hour content-discovery run that dies to a dropped VPN picks up where it stopped rather than re-walking the wordlist. Detect-It-Easy's `--heuristicscan --verbose` engine adds entry-point emulation, MSIL obfuscation modeling and a known-plaintext key-recovery pass, so a packed PE that matches no signature still yields a family guess before you spend a sandbox slot on it.
>
> dirsearch

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Exploitation & C2

Bugcrowd

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Bugcrowd is a crowdsourced security platform connecting organizations with ethical hackers to identify vulnerabilities through bug bounty programs.

> Bugcrowd added new controls for managing Request a Response interactions and introduced exportable Vulnerability Detection Reports for scan findings.

└──▷ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 **Vulnerability Detection Report generation** NEW 55

Adds Vulnerability Detection Report generation under Vulnerability Scanning, providing exportable scan findings for customer programs; accessible via Vulnerability Scanning > Generating Vulnerability Detection Report, where a completed scan can be selected and the report downloaded.

Review and export a Vulnerability Detection Report after running a vulnerability scan against your program.

> 📍 In the Bugcrowd platform, go to Vulnerability Scanning › Generating Vulnerability Detection Report, then select your scan and download the Vulnerability Detection Report.

– Includes exact navigation steps to generate the report. changelog-20260831-b4baa138

02 **Request a Response expiration and reply controls** NEW 35

Customers can now remove the expiration on a Request a Response, extending researcher response windows beyond the default deadline, and can respond directly to a Request a Response from within the platform.

– Names two controls but no UI path or mechanism detail. changelog-20260831-b4baa138

WAS THIS USEFUL? 👍 👎

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-08-31

NOTES ↗

width=device-width,initial-scale=1" name="viewport"/> <meta content=

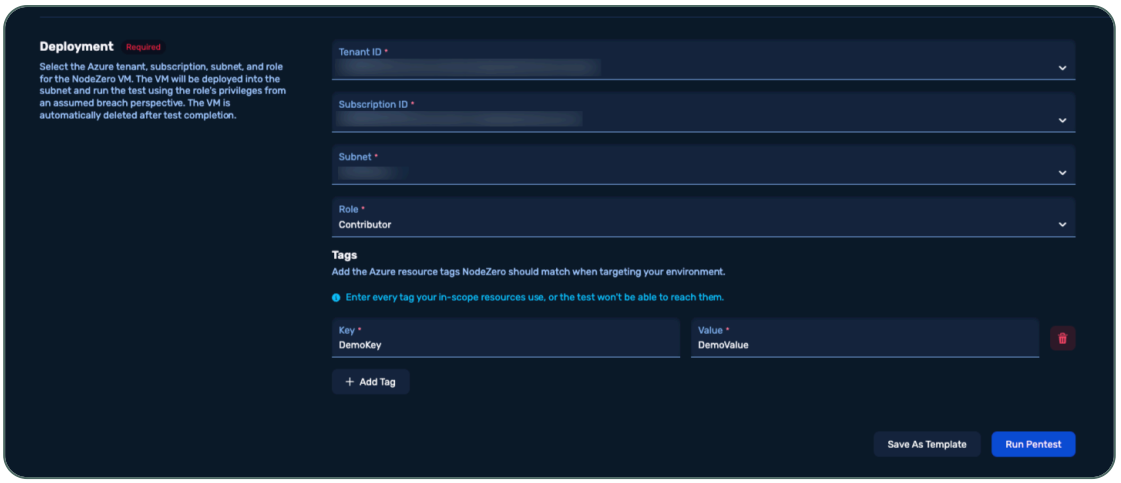
> NodeZero's biggest addition this window is Azure Cloud Pentest, a managed-identity-based assumed-breach test type paired with a new Cloud Connections flow for authorizing Azure tenants and subscriptions; it also picked up a WebApp test type and an AD Password Audit for Entra ID.

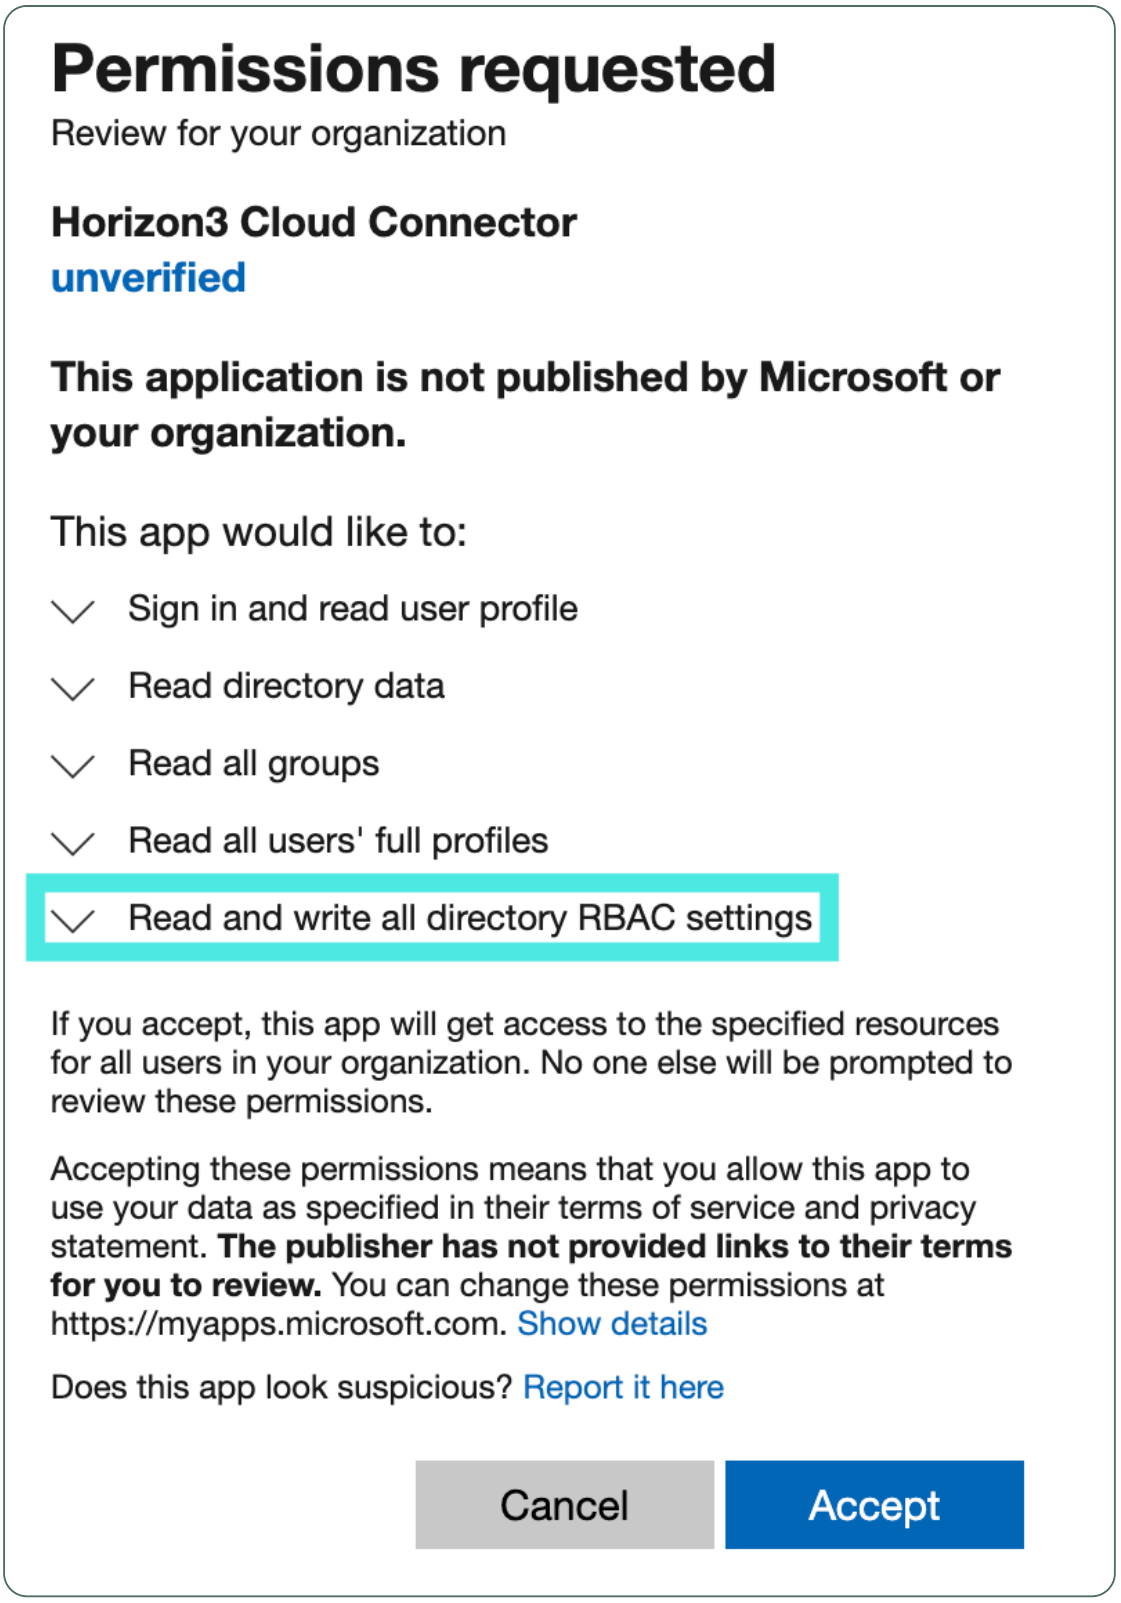
└──▷ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Azure Cloud Pentest test type and execution mechanics NEW

88

New Azure Cloud Pentest type, launched from Pentests > Run Assessment > Infrastructure > Azure Cloud Pentest and scoped per Tenant ID and Subscription ID, combines Azure API enumeration with an Internal Pentest in a single run — covering cloud resources and workloads on Azure VMs with no on-premises infrastructure required. NodeZero deploys a temporary `Standard_D4s_v3` VM into a subscriber-selected subnet and authenticates via a managed identity: a Horizon3 Cloud Connector Custom Role with 28 permissions across `Microsoft.Authorization`, `Microsoft.Compute`, `Microsoft.ManagedIdentity`, `Microsoft.Network` and `Microsoft.Resources` is provisioned, and dynamic just-in-time role assignment grants Global Reader at the Entra ID tenant level plus a subscriber-chosen subscription-level role (built-in Reader, Contributor, Virtual Machine Contributor or a custom role) only for the pentest's duration, revoking both at teardown. Pre-launch validation checks subnet IP space, NSG outbound connectivity on TCP 443, regional vCPU quotas (minimum 4 Total Regional vCPUs and 4 Standard DSv3 Family vCPUs), VM and network interface quotas, allowed VM SKUs, allowed locations, denied resource types (`Microsoft.Compute/virtualMachines`), and managed-disk policy, with a Tags field to satisfy required tag-value pairs or tag keys. Resources are deployed into a resource group named `H3-NodeZero-rg-[suffix]`, and all four resources (VM, OS disk, network interface, managed identity) are automatically removed when the pentest completes; a Real-Time View shows live progress and vulnerability findings as the test runs.


Deployment
Required
Select the Azure tenant, subscription, subnet, and role for the NodeZero VM. The VM will be deployed into the subnet and run the test using the role's privileges from an assumed breach perspective. The VM is automatically deleted after test completion.
Tenant ID *
Subscription ID *
Subnet *
Role *
Contributor
Tags
Add the Azure resource tags NodeZero should match when targeting your environment.
Enter every tag your in-scope resources use, or the test won't be able to reach them.
Key *
DemoKey
Value *
DemoValue
+ Add Tag
Save As Template
Run Pentest


Permissions requested
Review for your organization
Horizon3 Cloud Connector
unverified
This application is not published by Microsoft or your organization.
This app would like to:
Sign in and read user profile
Read directory data
Read all groups
Read all users' full profiles
Read and write all directory RBAC settings
If you accept, this app will get access to the specified resources for all users in your organization. No one else will be prompted to review these permissions.
Accepting these permissions means that you allow this app to use your data as specified in their terms of service and privacy statement. The publisher has not provided links to their terms for you to review. You can change these permissions at https://myapps.microsoft.com. Show details
Does this app look suspicious? Report it here
Cancel
Accept

Horizon3 Cloud Connector Custom Role

CustomRole

Showing 28 of 28 permissions

Type	Permissions	Description
Microsoft.Authorization		
Read	Get policy assignment	Get information about a policy assignment.
Read	Get policy definition	Get information about a policy definition.
Read	Get policy set definition	Get information about a policy set definition.
Read	Get role assignment	Get information about a role assignment.
Write	Create role assignment	Create a role assignment at the specified scope.
Delete	Delete role assignment	Delete a role assignment at the specified scope.
Microsoft.Compute		
Read	Get Disk	Get the properties of a Disk
Write	Create or Update Disk	Creates a new Disk or updates an existing one
Delete	Delete Disk	Deletes the Disk
Read	Get Virtual Machine	Get the properties of a virtual machine
Write	Create or Update Virtual Machine	Creates a new virtual machine or updates an existing virtual machine
Delete	Delete Virtual Machine	Deletes the virtual machine
Other	Power Off Virtual Machine	Powers off the virtual machine. Note that the virtual machine will continue to be billed.
Microsoft.ManagedIdentity		
Other	Assign User Assigned Identity	RBAC action for assigning an existing user assigned identity to a resource
Delete	Delete User Assigned Identity	Deletes an existing user assigned identity
Read	Get User Assigned Identity	Gets an existing user assigned identity
Write	Create/Update User Assigned Identity	Creates a new user assigned identity or updates the tags associated with an existing user assigned identity
Microsoft.Network		
Read	Get Network Interface	Gets a network interface definition.
Write	Create or Update Network Interface	Creates a network interface or updates an existing network interface.
Other	Join Virtual Machine to a network interface.	Joins a Virtual Machine to a network interface. Not Alertable.
Delete	Delete Network Interface	Deletes a network interface
Read	Get Network Security Group	Gets a network security group definition
Read	Get Virtual Network	Get the virtual network definition
Read	Get Virtual Network Subnet	Gets a virtual network subnet definition
Other	Join Virtual Network.	Joins a virtual network. Not Alertable.
Microsoft.Resources		
Read	Get Resource Group	Gets or lists resource groups.
Write	Create Resource Group	Creates or updates a resource group.
Delete	Delete Resource Group	Deletes a resource group and all its resources.

Name ↑		Type
H3-NodeZero-nic-2k2e0c	···	Network Interface
H3-NodeZero-osdisk-2k2e0c	···	Disk
H3-NodeZero-umai-2k2e0c	···	Managed Identity
H3-NodeZero-vm-2k2e0c	···	Virtual machine

Global Reader | Assignments

Privileged Identity Management | Microsoft Entra roles

Add assignments | Settings | Refresh | Export | Got feedback?

- Manage
  - Assignments
  - Description
  - Role settings

Eligible assignments | Active assignments | Expired assignments

2k2

Name	Principal name	Type
Global Reader		
H3-NodeZero-umai-2k2e0c	-	Service principal

Run an identity- and configuration-focused Azure pentest using the Reader role to test with minimal permissions.

> 📍 In the NodeZero portal, go to Pentests › Run Assessment › Infrastructure › Azure Cloud Pentest. Enter your pentest name, select your Tenant ID and Subscription ID, choose a subnet for NodeZero VM deployment, then set Role to `Reader`. Submit the form — NodeZero will run pre-launch validation checks before launching.

Enable deeper gray-box testing including infrastructure provisioning by selecting the Contributor role for a subscription with required tags.

> 📍 In the NodeZero portal, go to Pentests › Run Assessment › Infrastructure › Azure Cloud Pentest. Select Tenant ID and Subscription ID, pick a subnet, set Role to `Contributor`, and supply any required tag key-value pairs in the Tags field to satisfy subscription or resource group policy. Click Run Pentest after pre-launch validation passes.

Use a read-only IAM posture audit by assigning the `Reader` role at launch, so NodeZero enumerates Azure resources without write-level attack surface.

> 📍 In the NodeZero Portal, launch a new Azure Cloud Pentest › select your Cloud Connector (tenant) and target subscription › select the subnet with outbound internet access › set the subscription-level role to `Reader` › start the test.

Run deeper gray-box testing by assigning `Contributor` at launch, giving NodeZero write-level Azure API attack surface in addition to directory enumeration via `Global Reader`.

> 📍 In the NodeZero Portal, launch a new Azure Cloud Pentest › select your Cloud Connector (tenant) and target subscription › select the subnet with outbound internet access › set the subscription-level role to `Contributor` › start the test.

— Full mechanism, named roles, quotas and cleanup, but no CLI/API endpoint.

`snapshot-20260831`

---

02 **Cloud Connections setup for Azure tenants and subscriptions**

NEW 82

New `Pentests > Cloud Connections` section lets practitioners create and manage Azure cloud connections required to run Azure Cloud Pentests. Subscriptions can be added via direct Subscription ID entry or by uploading a CSV file with a required `id` column header. A tenant verification step reads Microsoft-verified domains from Microsoft Entra to confirm tenant ownership, and a permissions verification step confirms the `Horizon3 Cloud Connector Custom Role` PowerShell script ran successfully and the managed identity can authenticate. Once verified, practitioners can launch a test

directly via a `Run Pentest` button or add more tenants with `Add Another Tenant`, and a `Check Again` button recovers from admin consent polling timeouts without restarting the flow.

Export all Azure subscription IDs as a CSV to bulk-upload during connection setup, avoiding manual ID entry for large environments.

> 📍 ResourceContainers | where type == `microsoft.resources/subscriptions` | project id = subscriptionId

Navigate to the Azure Cloud Connection creation form in NodeZero Portal to start connecting your Azure tenant.

> 📍 In the NodeZero Portal, go to Pentests › Cloud Connections, click `+ Connection`, select `Azure`, fill in Tenant Name and Tenant ID, add subscriptions (or upload a CSV with an `id` column), click `Create Connection`, then accept the Horizon3 Cloud Connector consent prompt in Azure as a Global Administrator.

— Exact menu path, CSV column name and button labels given. `snapshot-20260831`

THINNER COVERAGE BELOW

03 **AD Password Audit for Azure Entra ID** NEW 39

New AD Password Audit feature audits Active Directory passwords within Azure Entra ID, accessed via Azure Entra ID > AD Password Audit after connecting an Azure tenant.

Audit Active Directory passwords in your Azure Entra ID environment to identify weak or compromised credentials.

> 📍 In the NodeZero console, navigate to Azure Entra ID › AD Password Audit, connect your Azure tenant, and launch the audit.

— Names the feature and navigation path but no mechanism detail. `snapshot-20260831`

04 **WebApp test type** NEW 18

Adds a new `WebApp` test type to NodeZero, enabling web application-focused penetration tests alongside existing test types.

— One-line mention with no mechanism or usage detail. `product docs`

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

ProjectDiscovery httpx-toolkit

Feedback ⓘ

1 RELEASE · 2026-08-31 NOTES ↗

The httpx-toolkit probes HTTP services in parallel to identify live hosts, technologies, status codes, titles, and other web metadata.

> httpx v1.11.0 makes page classification opt-in via a new knowledge-base flag and expands TLS fingerprint spoofing with Chrome and JA3 strategies.

└──▷ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Opt-in knowledge-base page classification IMPROVED

80

The `-kb` / `-knowledge-base` flag makes httpx's page-type classification (tagging pages as login, captcha, parked, etc.) opt-in rather than always-on, and results can be filtered with existing flags like `-fpt` (e.g. `-fpt parked,captcha`).

Enable knowledge-base page classification to tag pages as login, captcha, parked, etc., then filter on the result.

```
$ httpx -l targets.txt -kb -fpt parked,captcha -o live.txt
```

— Names exact flag, values, and a runnable command example. v1.11.0

THINNER COVERAGE BELOW

02 Chrome and JA3 TLS impersonation strategies NEW

52

The existing `-tls-impersonate` flag gains `chrome` and `JA3` strategies for more realistic TLS fingerprint spoofing.

— Names flag and new strategy values but gives no usage example. v1.11.0

WAS THIS USEFUL? 👍 👎

Nmap

Feedback ⓘ changes since 2026-08-27 CODE ↗

Nmap - the Network Mapper. Github mirror of official SVN repository.

Nmap adds separate IPv4/IPv6 trie structures and a new SAP HANA database service probe

└──▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-08-27
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-27
```

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian Self-Hosted

Feedback ⓘ 1 RELEASE · 2026-08-31

CHANGELOG ↗

GitGuardian Self-Hosted detects and prevents secrets leakage across your organization's Git repositories on your own infrastructure.

> GitGuardian Self-Hosted 2026.8.0 adds an in-cluster ClickHouse analytics database and a beta Remote MCP server, alongside a self-service custom validity API, AI-powered incident search, new public API token management, and over 25 new or improved secret detectors and checkers.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 In-cluster ClickHouse analytics database NEW

78

Adds `clickhouse-backup` sidecar and ClickHouse as an optional in-cluster column-oriented database, backed by a StatefulSet and S3-compatible object storage, deployable via Helm or KOTS to power Endpoint Protection analytics. ClickHouse is optional in 2026.8 but becomes mandatory in 2027.1.0, requiring S3-compatible object storage (Azure Blob Storage and GCS support in progress) to be provisioned before that upgrade.

— Names deployment mechanism and future mandatory requirement, no command shown

changelog-20260831-63685e46

02 Unified filter bar and AI-powered incident search NEW

78

Enables the unified filter bar on self-hosted, supporting incident search by secret value, commit author, or file path. Plain-English search (AI Filters) activates once a workspace administrator enables the External LLM toggle under `Settings > Workspace > AI` and configures a reachable LLM provider, after which AI Filters appears in the Incidents page filter bar.

Enable Plain-English AI Filters for incident search on a self-hosted workspace after connecting an LLM provider.

> 📍 In the dashboard, go to Settings › Workspace › AI and enable the External LLM toggle, then configure a reachable LLM provider. The AI Filters option will appear in the unified filter bar on the Incidents page.

— Gives exact settings path and activation steps changelog-20260831-63685e46

03 GitGuardian Remote MCP server NEW

68

Adds a GitGuardian Remote MCP server deployment (Helm only, beta, disabled by default) that runs in-cluster so no traffic or secrets leave the network. Point Claude, Cursor, or any MCP client at the instance to scan for leaked credentials, manage incidents, and deploy honeytokens from the IDE.

— Explains scope and clients but no setup command given `changelog-20260831-63685e46`

---

04 **Custom validity status via API** NEW 66

Adds a Set your own validity status API call that accepts `valid` or `invalid` for secrets on internal systems GitGuardian cannot reach. Affected incidents display Custom validity with the date set, recorded in both the incident activity log and the audit log.

— Names accepted values and log surfaces, no endpoint path given

`changelog-20260831-63685e46`

---

05 **New and upgraded secret detectors and checkers (v2.167–v2.169.1)**

NEW 62

Adds detectors across v2.167–v2.169.1: Azure SQL Credentials, Azure Machine Learning Key, Azure Video Indexer Key, Azure Health Insights Key, Azure Fluid Relay Key, LiteLLM API Key, DashScope API Key, Apify Proxy Token, Authress Access Key, two Gemfury token variants, Azure Immersive Reader Key, DashScope Coding Plan API Key, DashScope Token Plan API Key, Prefect Cloud API Key, Slack Webhook Trigger URL, n8n API Key, Atlassian Auth Key, AWS Cognito Access Token, OpenAI Access Token, Vercel OIDC Token, OpenCode API Key, Lovable Access Token, and Lovable Git Token. Adds or upgrades checkers: LiteLLM Without Host, LaunchDarkly SDK Key, Llama Cloud API Key, New Relic APM License Key, npm Token, Hume AI API Key, Alibaba Cloud IAM, GitLab Deploy Token, and Supabase Service Role JWT.

— Names every detector and checker, no direct user action `changelog-20260831-63685e46`

THINNER COVERAGE BELOW

---

06 **Azure DevOps ggshield pipeline extension** NEW 55

Adds a new Azure DevOps extension that inserts a `ggshield` scan job into every pipeline in an organization via a single pipeline decorator.

— Names the mechanism and tool but no install steps `changelog-20260831-63685e46`

---

07 **Automatic monitoring and scan toggles for collaboration sources**

NEW 55

Adds automatic monitoring and automatic scan toggles for Jira Cloud, Jira Data Center, Confluence Cloud, Confluence Data Center, and Microsoft Teams sources.

— Names all five sources but not the toggle location `changelog-20260831-63685e46`

08 **Detector validators option on custom sources** NEW 50

Adds the Only use detectors with validators option as settable on custom sources through the public API.

– Names exact config field but not the endpoint changelog-20260831-63685e46

09 **GitHub check runs page in Perimeter** NEW 50

Adds a dedicated GitHub check runs page in the Perimeter section showing status, repository, commit, and linked pull request, with in-dashboard re-run or skip actions.

– Clear UI location named, fields and actions listed changelog-20260831-63685e46

10 **HTML/Markdown scanning extended to more token types** IMPROVED 50

Extends HTML/Markdown scanning in v2.168 to npm, GitHub Personal Access, GitHub OAuth, GitHub App, and Slack Application tokens.

– Names token types covered, no scan trigger detail changelog-20260831-63685e46

11 **Personal Access Token and Service Account Token creation via public API**

NEW 45

Adds Personal Access Token and Service Account Token creation through the public API for automated rotation workflows.

– Names token types and use case but no endpoint or example changelog-20260831-63685e46

12 **Tunable emptyDir limits for ml-secret-engine** IMPROVED 40

Adds tunable `emptyDir` volume limits for the `ml-secret-engine` pod in the Helm chart.

– Names the exact Helm chart setting, no values given changelog-20260831-63685e46

13 **Fine-grained GitHub PAT for private members monitoring** NEW 40

Adds fine-grained GitHub personal access token support for private members monitoring in the Public Monitoring perimeter.

– Names feature area but no setup steps changelog-20260831-63685e46

14 **Detector removals for discontinued services** BREAKING 40

The Grafbase Access Token detector is removed and the Amadeus OAuth Credentials detector is disabled, as both services were discontinued.

— Names both detectors and reason, no remediation steps changelog-20260831-63685e46

15 **@-mentions in incident notes and feedback** NEW 35

Adds `@ -mention` support for members and teams in incident notes and feedback, triggering email notifications on both Internal and Public Monitoring incidents.

— Bare description of behaviour, no config or scope detail changelog-20260831-63685e46

16 **Display as Tree view removed from source pages** BREAKING 30

The Display as Tree view is removed from integration source pages.

— Names removed feature but no migration guidance changelog-20260831-63685e46

17 **Audit log event for PAT-triggered scans** IMPROVED 25

Scans triggered via a Personal Access Token now emit an audit log event.

— Bare statement, no detail on log fields or access changelog-20260831-63685e46

└──▷ ALSO FROM THESE RELEASES

.
.
.

└──▷ BREAKING ON UPGRADE

! The Display as Tree view is removed from integration source pages.

! ClickHouse is optional in 2026.8 but becomes mandatory in 2027.1.0, requiring object storage (S3-compatible; Azure Blob Storage and GCS in progress) to be provisioned before upgrading to that release.

! The Grafbase Access Token detector is removed and the Amadeus OAuth Credentials detector is disabled, as both services were discontinued.

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS

Feedback ⓘ 2 RELEASES · 2026-08-31 BLOG ↗

GitGuardian SaaS is a cloud-based platform that detects and prevents secrets and sensitive data from being exposed in Git repositories.

> GitGuardian's biggest addition this window is AWS S3 bucket scanning that extends secret detection beyond Git into cloud storage (including compressed archives), alongside AI-agent verdicts on public leaks that replace the old ML risk score, new auto-monitoring for Jira, Confluence and Microsoft Teams, GitHub source health handling, and SSO/SCIM name derivation.

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

01 **AWS S3 bucket scanning integration** NEW 90

Adds AWS S3 to GitGuardian's monitored perimeter, running historical scans of accumulated objects — including backups, CI/CD logs, and Terraform state files — to surface exposed credentials as incidents, and extends scanning into supported `ZIP` and `tar.gz` archives so secrets hidden inside compressed objects are found. Connects via an OIDC-assumed read-only IAM role (no long-lived keys), lets you select individual buckets and assign them to teams, and routes findings into existing incident workflows with the usual severity, validity, playbooks, and notifications. Scanning defaults to the integration's region and the S3 Standard storage class; cross-region scanning and other storage classes are opt-in to control retrieval and transfer costs.

⋮

When onboarding S3 to your monitored perimeter to surface years of credential debt in backup, CI/CD, and infrastructure buckets.

> 📍 In GitGuardian, go to Perimeter › Sources › AWS S3, select the buckets you want monitored, and trigger a historical scan to analyze existing objects for exposed credentials.

Connect an AWS S3 bucket for secret scanning without handing over long-lived keys — set up the OIDC provider and scoped IAM role, then register the account in GitGuardian.

> 📍 In the GitGuardian dashboard, go to Integrations › AWS S3 and follow the three-step setup: (1) create the OIDC provider in AWS, (2) create a read-only IAM role scoped to the buckets you want monitored, (3) enter your AWS account ID and role name in the integration settings. Then select individual buckets, assign them to teams, and trigger a historical scan.

— Names OIDC/IAM setup, archive formats, storage class defaults; steps but no CLI command

changelog-20260831-6236c8cf launch-20260831-00fa89ef

02 **AI-agent verdicts on public secret leaks** BREAKING 80

Adds AI-agent analysis of public secret incidents, producing a Company-related verdict (Related , Uncertain , or Unrelated) exposed as a column, a filter, and three ready-made saved views, plus a new Analysis tab showing the company identified, the agents' reasoning, the risk score with rationale, and which agent concluded. This replaces the ML risk score on public incidents and supersedes the Context related to company tag; risk score filtering is no longer available for notifications on public incidents.

Triage public leaks by company relevance — filter to only incidents the AI agents flagged as Related to your company so your team works the leaks most likely to be yours first.

> 📍 In the GitGuardian dashboard, open Public Monitoring › Incidents, then use the Company-related filter and select Related . Alternatively, open one of the three ready-made saved views (Related , Uncertain , Unrelated) from the saved views menu.

— Explains verdict mechanism and named UI surfaces, but no API/flag

changelog-20260831-6236c8cf

03 **Auto-monitoring for Jira, Confluence and Teams** NEW 60

Adds automatic monitoring and automatic scan toggles for Jira Cloud , Jira Data Center , Confluence Cloud , Confluence Data Center , and Microsoft Teams — new spaces, projects, and channels are picked up without manual addition.

— Names every covered source but no config/UI path given changelog-20260831-6236c8cf

THINNER COVERAGE BELOW

04 **Source health management for GitHub** IMPROVED 50

Sources health management for GitHub now pauses real-time ingestion and historical scans on unreachable GitHub sources, auto-resumes scanning on recovery, and surfaces an actionable recovery step.

— Describes mechanism but no UI location or example given changelog-20260831-6236c8cf

05 **SSO/SCIM sign-in without name attributes** IMPROVED 40

Adds SSO/SCIM sign-in support when the identity provider does not supply first and last name, deriving these values from the user's email address instead.

— Clear behaviour but no config key or setup steps named changelog-20260831-6236c8cf

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! When AI-agent analysis is enabled, the agents' risk score replaces the ML risk score on public incidents, and the `Context related to company` tag is superseded by the Company-related verdict; risk score filtering is no longer available for notifications on public incidents.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Socket

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Socket detects supply chain attacks by analyzing open source package dependencies for malicious code, vulnerabilities, and suspicious behaviors.

> Socket rolled out an Asana ticketing integration for alerts, promoted PHP/Composer support to Beta with GA reachability analysis for Team plans, and added experimental Firefox extension scanning for malicious add-on behavior.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Asana integration for alert ticketing** NEW 65

Teams can create and manage Asana tasks directly from Socket alerts, either manually or via ticketing rules that filter by event, type, priority, and repository, with configurable sync between task and alert status. Available in beta on Business and Enterprise plans.

– Describes rule filters and sync but no exact UI path or API changelog-20260831-4b8949dc

02 **Firefox extension scanning** NEW 60

New scanning covers every extension in Mozilla's official add-ons directory, monitoring for permission changes, network activity, malware, credential and clipboard theft, data exfiltration, remote loading, obfuscation, and impersonation. Available in Experimental for Enterprise customers.

– Lists detection categories but no setup steps or endpoint changelog-20260831-4b8949dc

THINNER COVERAGE BELOW

03 **PHP and Composer support promoted to Beta** IMPROVED 50

PHP and Composer support moved from Experimental to Beta and is now available to all Socket users; PHP reachability analysis is generally available for Team plan customers and above without requiring Experimental access.

– States plan tiers and status change but no config or command changelog-20260831-4b8949dc

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback

1 RELEASE · 2026-08-31

CHANGELOG ↗

A fully-managed, enterprise-scale solution to control, secure, and distribute software packages and containers.

> Cloudsmith added native Nix package hosting, extended cooldown policies to Conda and Docker packages, introduced a Libraries tab for vendor-curated upstreams, and brought usage tracking into the Broadcasts app.

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

01 **Nix package format support** NEW 70

Cloudsmith now supports the Nix package format, enabling hosting of public or private Nix binary caches with Ed25519-signed packages, along with upstream proxying and caching against NixOS release channels.

— Names mechanism (Ed25519 signing, proxying) but no config keys or commands.

changelog-20260831-681ffc66

THINNER COVERAGE BELOW

02 **Cooldown policies for Conda and Docker** IMPROVED 55

Cooldown policies, which protect supply chains from newly published versions that may carry malware or lack sufficient community scrutiny, are now extended to cover Conda and Docker packages in addition to prior formats.

— Explains purpose and scope but no thresholds or config surface named.

changelog-20260831-681ffc66

03 **Libraries tab for upstream registries** NEW 50

A new Libraries tab lets users quickly configure library registries as upstreams, giving access to vendor-curated, pre-vetted, and security-scanned open source package feeds.

— Names UI location as starting point but lacks further mechanism detail.

changelog-20260831-681ffc66

04 **Broadcasts consumption tracking** IMPROVED 45

The Broadcasts app now adds download tracking, entitlement token usage, and client log visibility directly within the app.

— Lists three additions but no config keys or endpoints given.

changelog-20260831-681ffc66

WAS THIS USEFUL?  

GitHub

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

GitHub is a web platform for version control and collaborative software development using Git repositories.

> GitHub's August 2026 Copilot in VS Code update centers on cross-window Agent Host sessions, a shareable-context `/btw` side chat and `/rubber-duck` second-opinion command, model switching between Anthropic and Copilot subscriptions, multi-language on-device dictation, and a batch of chat/Agents-window layout and integrated-browser additions.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 Integrated browser comment and auto-reload features NEW

80

Adds comment-on-elements capability in the integrated browser to select and annotate HTML elements for targeted UI feedback on batch agent edits; adds automatic reload of local HTML files when they change on disk; and supports setting the integrated browser as the default editor for HTML files via the editor associations setting (Editor > Associations, mapping `*.html`).

Set the integrated browser as the default editor for HTML files so local previews open automatically instead of the text editor.

> 📍 In VS Code Settings, go to Editor › Associations and add an entry mapping *.html to the integrated browser editor.

— Exact settings path and mapping given as example.

GitHub Copilot in VS Code, August 2026 …

02 /rubber-duck complementary model command NEW

75

Adds an experimental `/rubber-duck` command usable in a Copilot Agent Host session to ask a complementary model to surface missed details and edge cases before merging.

Get a second opinion from a complementary model on the current agent session to catch missed edge cases before merging.

```
$ /rubber-duck
```

— Exact command with usage example, mechanism thin.

GitHub Copilot in VS Code, August 2026 …

03 Side chat via /btw command NEW

70

Adds a `/btw` slash command that opens a side chat sharing the primary chat's context and prompt cache, while the main chat continues running.

— Names exact command and mechanism, no further usage detail.

GitHub Copilot in VS Code, August 2026 …

04 **Multi-language dictation support** NEW 65

Adds multi-language dictation using an on-device model with automatic language detection, user-level and workspace-level dictation instructions for custom transcript cleanup aligned to project terminology, and shell-aware cleanup that preserves command syntax when dictating in the terminal instead of inserting spoken punctuation literally.

— Describes mechanism and scope but no config key named.

GitHub Copilot in VS Code, August 2026 …

05 **Agents window and chat layout options** IMPROVED 62

Adds chat sticky scroll to keep the current prompt pinned while scrolling through long conversations; supports arranging chats side by side in horizontal or vertical groups with layout restored on return to the session; adds a default single-pane layout that opens session details and editors in one side pane with adaptive-width diffs; enables switching between regular and diff editors from the breadcrumb bar in the Agents window; and allows terminal output expanded in chat to reflow as the view is resized.

— Multiple UI locations named but no shortcuts or settings keys.

GitHub Copilot in VS Code, August 2026 …

THINNER COVERAGE BELOW

06 **Full-transcript chat search** NEW 55

Adds text search across the complete chat transcript with options for case matching, whole words, and regular expressions.

— Describes options but no UI path or shortcut given.

GitHub Copilot in VS Code, August 2026 …

07 **Token usage hover in chat responses** NEW 55

Shows input, cached input, and output token usage per model in a chat turn by hovering over the response footer.

— Exact hover location given, no further breakdown.

GitHub Copilot in VS Code, August 2026 …

08 **Continue Copilot or Claude sessions across apps** NEW 53

Adds ability to view and continue recent Copilot or Claude agent sessions created in other applications from the Sessions list in VS Code.

— Names Sessions list location, little further detail.

GitHub Copilot in VS Code, August 2026 …

09 **Agent Plugins 1.0 standard support** NEW 50

Supports the Agent Plugins 1.0 standard, enabling portable agent plugin installation across VS Code and compatible agent clients.

— No install command or config surface shown. GitHub Copilot in VS Code, August 2026 …

10 **Prompt timeline in transcript gutter** NEW 50

Adds a prompt timeline control in the transcript gutter to navigate to specific prompts and review related file changes.

— UI location named but no exact steps. GitHub Copilot in VS Code, August 2026 …

11 **Switch between Anthropic and Copilot models in Claude sessions**

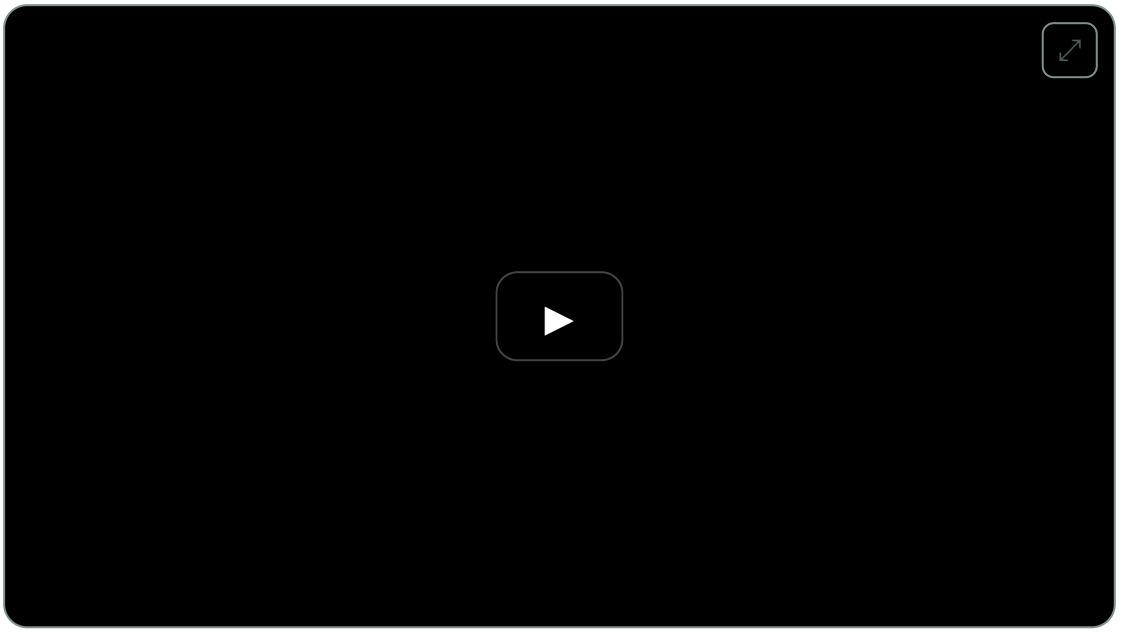
NEW 48

Enables switching between models from an Anthropic subscription and a Copilot subscription at any time within Claude sessions.

— Describes capability but no menu path given. GitHub Copilot in VS Code, August 2026 …

12 **Cross-window agent sessions via Agent Host** NEW 40

Enables connecting multiple VS Code windows to the same agent session via Agent Host.



— Brief description, no setup steps given. GitHub Copilot in VS Code, August 2026 …

13 **Sign-in-free Agents window with Claude API key** NEW

40

Adds an experimental setting to open the Agents window without GitHub sign-in when Claude is configured with an API key.

— Setting name not specified, limits actionability.

GitHub Copilot in VS Code, August 2026 …

14 **Hybrid Markdown/diff editor** NEW

35

Adds an experimental hybrid Markdown editor combining diff viewing and Markdown editing in a single experience.

— Thin description, no usage path given. GitHub Copilot in VS Code, August 2026 …

└──▷ ALSO FROM THESE RELEASES







WAS THIS USEFUL? 👍 👎

JFrog

Feedback ⓘ 2 RELEASES · 2026-08-24 BLOG ↗

JFrog is a DevOps platform providing software supply chain management, artifact management, and continuous integration/delivery automation tools.

> JFrog introduced Platform Federation, a peer-to-peer mesh that keeps projects, access control, curation policies, and repositories synchronized across every self-managed and SaaS site in a software supply chain without manual per-site work.

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

01 Curation Federation for blocking rules and waivers NEW

75

Curation Federation propagates open-source blocking rules, waivers, and approvals from any peer site to every regional pipeline worldwide within seconds of being set. To use it, go to Curation > Policies in the JFrog Platform UI, create or update a blocking rule or waiver on any peer site, and Platform Federation propagates it automatically to every connected site in the mesh.

When a new zero-day drops, set a curation blocking rule once on any peer site and have it enforced across every regional pipeline within seconds — no per-site rollout scripts needed.

> 📍 In the JFrog Platform UI, go to Curation › Policies, create or update a blocking rule or waiver on any peer site, and Platform Federation propagates it to every connected site in the mesh automatically.

When your security team publishes a zero-day blocking rule and needs it enforced globally within seconds rather than rolling it out region by region.

> 📍 In the JFrog Platform, go to Curation › Policies, create or update a blocking rule or waiver on any peer site, and Platform Federation propagates it to every connected regional pipeline automatically.

— Names exact UI path and describes propagation mechanism and timing.

Introducing JFrog Platform Federation: … launch-20260824-c2b73561

02 Automated Repository Provisioning via Stage-to-Site mapping

NEW

75

Automated Repository Provisioning creates and configures repositories on the correct sites automatically based on a Stage-to-Site mapping, eliminating per-repo and per-site manual administration. To use it, go to Platform Federation > Stage-to-Site Mapping in

the JFrog Platform UI, define which release stages map to which regional sites, then add the new site to the mesh so required repositories are provisioned automatically.

When onboarding a new region, define your Stage-to-Site mapping once so Automated Repository Provisioning creates the correct repositories on the new site without tickets or manual REST API calls.

> 📍 In the JFrog Platform UI, go to Platform Federation › Stage-to-Site Mapping, define which release stages map to which regional sites, then add the new site to the mesh — required repositories are provisioned automatically.

When onboarding a new region so that all required repositories appear automatically without filing tickets or running scripts for each repo.

> 📍 In the JFrog Platform, go to Platform Federation › Stage-to-Site Mapping, define which release stages are assigned to the new site, and Automated Repository Provisioning replicates the required repositories to that site.

— Names exact UI path and mapping mechanism for repo creation.

`Introducing JFrog Platform Federation: …` `launch-20260824-c2b73561`

---

03 **Platform Federation peer-to-peer mesh architecture** NEW

60

JFrog Platform Federation introduces a peer-to-peer mesh architecture for distributed JFrog sites, replacing hub-and-spoke and script-driven sync with autonomous, continuous background synchronization of projects, access control, policies, artifacts, and repositories. It supports both self-managed and JFrog Cloud SaaS sites — including airgapped and private-cloud deployments — in the same federated mesh, enabling phased SaaS migrations without breaking centralized security or identity control, and provides a Federated Topology View showing the global peer mesh across regions with entities and policies synchronized in real time.





— Explains architecture and deployment scope but gives no config or command surface.

Introducing JFrog Platform Federation: … launch-20260824-c2b73561

THINNER COVERAGE BELOW

04 Projects Federation for workspaces and access control NEW

53

Projects Federation synchronizes multi-tenant workspaces, project- and global-level roles, SDLC release-stage gates, and identity access tokens across every peer site in the mesh in seconds.

— Names what syncs but no UI path or config key given.

Introducing JFrog Platform Federation: … launch-20260824-c2b73561

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Bright Security

Feedback ⓘ 1 RELEASE · SEEN 2026-08-31 PRESS ↗

Bright is a dynamic application & API security testing platform that offers automation, fast scans, minimal false positives, and comprehensive security testing.

> Bright Security added active detection for EXIF metadata leakage in image upload flows, validating the full processing path and producing reproducible proof artifacts.

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

01 EXIF metadata leakage detection for image uploads NEW

83

A new check actively tests multipart and JSON upload flows using JPEG and PNG canary images embedded with controlled metadata in fields such as `UserComment`, `ImageDescription`, `Artist`, timestamp fields, and GPS coordinates; a finding is created only when a metadata field matches exactly or a GPS coordinate pair is confirmed preserved in the served image. The scanner validates the full image processing path — upload endpoint, application response, image discovery via response body, redirect headers, JSON/HTML responses, image links, and rendered DOM — rather than flagging storage configuration alone, and produces evidence including the original upload request, the application response, the retained EXIF fields, and the URL where the image was found for use as a testable AppSec control in CI/CD.

— Detailed mechanism and named fields, but no CLI/API to invoke it.

launch-20260831-830cadf9

WAS THIS USEFUL? 👍 👎

Dalfox

Feedback

1 RELEASE · 2026-08-31 NOTES ↗ CODE ↗

Dalfox is a powerful open-source XSS scanner and utility focused on automation.

> Dalfox v3.2.2 adds a shell completion subcommand, a canonical JavaScript payload selector, new WAF fingerprints, and fixes CLI flag precedence over config files.

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

01 **JavaScript payload selector for XSS payloads** NEW 75

Adds `dalfox payload javascript` selector to print canonical JavaScript execution payloads used in JS-string contexts.

Generate a list of canonical JavaScript XSS payloads to review or pipe into a custom scan.

```
$ dalfox payload javascript
```

— Runnable command with example provided v3.2.2

02 **Shell completion generation via `dalfox completion`** NEW 60

Adds the `dalfox completion` subcommand to generate shell completions, installed alongside the man page by packages.

— Names the subcommand but no example run shown v3.2.2

THINNER COVERAGE BELOW

03 **Slash-separated attribute breakout payloads** NEW 55

Adds slash-separated attribute breakout payloads to `dalfox payload` for bypassing space-stripping filters.

— Names payload type but no example command v3.2.2

04 **CLI flag precedence over config files fixed** IMPROVED 40

Config files no longer override flags explicitly typed on the CLI, fixing flag-precedence for all subcommands.

— Describes fix but no specific config key named v3.2.2

05 **WAF fingerprint detection additions** NEW 30

Adds WAF fingerprints for Wallarm, NAXSI, and SafeLine.

— Only named in summary, no mechanism detail v3.2.2

WAS THIS USEFUL? 👍 👎

Acunetix 360

Feedback ⓘ 1 RELEASE · 2026-08-27 BLOG ↗

Invicti Web + API is an end-to-end web security scanner that offers a 360 view of an organization's security. Allowing you to take control of the security of all you web applications, web services, and APIs to ensure long-term protection.

> Acunetix 360 shipped opt-in HAR file capture for authentication verification, letting practitioners inspect the full HTTP exchange during auth checks.

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

01 Opt-in HAR generation for authentication checks NEW

43

Adds opt-in HAR file generation for authentication verification, enabling practitioners to capture and inspect the full HTTP request/response exchange during auth checks.

— Names HAR format and purpose but no toggle location or config key given.

v26.8.1 – 27 August 2026 (2026-08-27)

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

> Elcomsoft System Recovery 8.38 adds a built-in BitLocker TPM exploit library and checksum logging for disk images captured during forensic triage.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 Built-in BitLocker TPM exploit library NEW 85

Elcomsoft System Recovery 8.38 adds a built-in BitLocker TPM exploit library supporting the `YellowKey` and `GreatXML` exploits. During installation, the installer offers to build a second USB flash drive loaded with the chosen exploit; the target machine is then booted with both drives attached to attack a locked drive.

To leverage the BitLocker TPM exploit against a locked drive, use the installer to build a second USB flash drive with the YellowKey or GreatXML exploit, then boot the target machine with both drives attached.

📍 During Elcomsoft System Recovery 8.38 installation, when prompted, select Create BitLocker exploit USB, choose either YellowKey or GreatXML, and write to a second USB flash drive.

— Names exploits and exact install steps, fully actionable changelog-20260831-c2359679

THINNER COVERAGE BELOW

02 Checksum logging for disk images NEW 55

Adds checksum logging for disk images captured during triage, producing a verifiable record of the disk image for chain-of-custody purposes. Enabled via the Create Disk Image imaging options screen before starting acquisition.

When imaging a suspect drive during triage, use the new checksum logging to produce a verifiable record of the disk image for chain-of-custody.

📍 During bootable triage in Elcomsoft System Recovery 8.38, select Create Disk Image, then enable checksum logging on the imaging options screen before starting the acquisition.

— Clear purpose and UI path, but no format/algorithm detail changelog-20260831-c2359679

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Wireblast

Feedback ⓘ 1 RELEASE · 2026-08-31 NOTES ↗

Wireblast is an AF_XDP traffic generator for measuring and replaying high-rate Linux network traffic.

> Wireblast v0.6.1 adds a new mlx5 Direct Verbs backend for receive-mode traffic generation and tightens `AF_XDP` safety checks around queue binding and VLAN filtering.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 **mlx5 Direct Verbs backend with --rx-mode** NEW 75

Adds a new mlx5 Direct Verbs backend, enabling `--rx-mode` receive support directly on Mellanox mlx5 NICs. For example: `sudo wireblast -i eth1 --mode receive --rx-mode udp-port --rx-port 9000 -d 60s`.

Use receive mode on an mlx5 Mellanox NIC now that the Direct Verbs backend supports `--rx-mode`.

```
$ sudo wireblast -i eth1 --mode receive --rx-mode udp-port --rx-port 9000 -d 60s
```

— Names backend, flag, and gives runnable example v0.6.1

THINNER COVERAGE BELOW

02 **AF_XDP receive safety checks** IMPROVED 40

Wireblast now warns when an `AF_XDP` receive binds fewer queues than the NIC actually has, alerting users to potential missed traffic, and refuses an `AF_XDP` receive outright when the kernel's VLAN filter would silently drop that traffic, preventing silent data loss.

— Describes behavior but names no flags or config v0.6.1

WAS THIS USEFUL? 👍 👎

◆ IDS / IPS

CrowdSec

Feedback ⓘ 1 RELEASE · 2026-08-31 NOTES ↗

CrowdSec - the open-source and participative security solution offering crowdsourced protection against malicious IPs and access to the most advanced real-world CTI.

> CrowdSec's v1.8.0 cycle centers on WAF bot detection with a challenge-and-fingerprinting mode and a wide set of AppSec observability metrics, alongside a new native Kubernetes log datasource, DNS lookup caching, and HTTP helpers in the expression language.

└──▷ **WHAT SHIPPED · 7 FEATURES** most completely described first ? what's the number?

01 **DNS lookup caching via `dns_cache`** NEW 85

Adds a `dns_cache` configuration block with `ttl` (default `1h`), `negative_ttl` (default `5m`), and `size` (default `16384`) keys, caching DNS lookups with an LRU cache to reduce repeated resolver calls. Example tuning for large or frequently-rotating address sets: `ttl: 30m`, `negative_ttl: 5m`, `size: 32768`.

Tune DNS caching when resolving a large, diverse set of addresses or when a vendor rotates DNS records frequently.

yaml

```yaml
dns_cache:
  ttl: 30m
  negative_ttl: 5m
  size: 32768
```

— Exact config keys, defaults, and a runnable YAML example. product docs

02 **AppSec challenge observability metrics** NEW 80

Adds `appsec-challenge` and `appsec-challenge-infra` metric categories to `cscli`, the latter exposing internal AppSec challenge runtime upkeep metrics: signing-key rotation, JS re-obfuscation, and cache eviction. Adds matching Prometheus counters `cs_appsec_challenge_requested_total`, `cs_appsec_challenge_submitted_total`, `cs_appsec_challenge_accepted_total` (labeled `kind=solved` for valid submissions or `kind=granted` via GrantChallengeCookie(reason)), `cs_appsec_challenge_rejected_total` (labeled `kind=protocol` for crypto-PoW failures, `kind=submission` via RejectSubmission(reason), or `kind=cookie` for invalid incoming cookies), `cs_appsec_challenge_exempt_total` for requests exempted via ExemptFromChallenge(reason) (verified bots like `gptbot`, well-known paths), and `cs_appsec_fingerprint_mismatch_total` labeled by `severity` from

EvaluateMismatches(). Runtime internals for CPU/memory troubleshooting are also exposed: `cs_appsec_challenge_kepoch_generated_total`, `cs_appsec_challenge_kepoch_evicted_total`, `cs_appsec_challenge_reobfuscation_total` (labeled `bundle`, with `bundle=dynamic` for the per-epoch signing module), and `cs_appsec_challenge_dynamic_module_evicted_total`.

— Every metric name and label given, ready to scrape. product docs

---

03 **WAF challenge-based bot detection** NEW 65

v1.8.0 adds WAF bot detection via `waf challenge mode`, which serves a challenge-and-fingerprinting page to clients before granting site access, with configurable score thresholds and cookie-disabled error messaging. `RejectSubmission` now generates an alert so WAF submission rejections surface in the alert stream, and a configurable auth timeout from WAF to LAPI was added.

— Mechanism described but no config key or command example given. v1.8.0

---

04 **Kubernetes log datasource** NEW 65

Adds a dedicated Kubernetes datasource that fetches logs directly from the k8s API server, streaming container logs from pods selected by labels, giving CrowdSec native k8s log acquisition without an agent proxy. Exposes a `cs_kubernetessource_hits_total` Prometheus counter labeled by `datasource_type`.

— Names the datasource and its metric but no config snippet. v1.8.0

---

05 **Machine heartbeat metrics** NEW 60

Adds `cs_machines_heartbeat_seconds` for tracking agent heartbeat timing and `cs_machines_last_heartbeat_timestamp`, a gauge exposing the Unix timestamp of the last heartbeat received from each machine.

— Names both metrics but no example usage given. v1.8.0

---

06 **cscli console status and hub list improvements** IMPROVED 60

Revamps `cscli console status` into an all-in-one status view for console enrollment and feature status, and improves `cscli hub list` output.

Check the all-in-one CrowdSec console enrollment and feature status after upgrading.

```
$ cscli console status
```

— Names both commands with a runnable check command. v1.8.0

THINNER COVERAGE BELOW

07 **HTTP helpers in expression language** NEW 35

Adds HTTP helpers to the expression language, enabling parsers and scenarios to query external services inline.

— No helper names or syntax given, just capability described. v1.8.0

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Google distroless

Feedback ⓘ changes since 2026-08-01

CODE ↗

Language focused docker images, minus the operating system.

Distroless hardens Node.js supply-chain trust with GPG signature verification and an `update-node-keys` subcommand in `knife`.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-01
https://github.com/GoogleContainerTools/distroless.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-01
```

└──▷ TRY IT

Generate an SPDX JSON SBOM for a generic (non-Debian) package such as a Node.js binary archive.

```
$ go run private/pkg/package_spdx.go \
  --package_name nodejs \
  --version 22.23.2 \
  --purl 'pkg:generic/nodejs@22.23.2?download_url=https://nodejs.org/'
\
  --sha256 <sha256-of-archive> \
  --url https://nodejs.org/dist/v22.23.2/node-v22.23.2-linux-
x64.tar.gz \
  --id nodejs-22.23.2 \
  --supplier 'Organization: OpenJS Foundation' \
  --output nodejs.spdx.json
```

Exposes `update-node-keys` subcommand in `knife` to sync committed Node.js release public keys from the upstream `nodejs/release-keys` repository, callable directly from CI workflows.

Verifies Node.js `SHASUMS256.txt.asc` via GPG `--decrypt` against committed release keys stored in `knife.d/nodejs_keys.asc`, replacing the previous unsigned plaintext fetch — no network access for key material.

Imports Node.js release public keys into an isolated temporary `GNUPGHOME` keyring (using `gpg --batch --import`) so GPG verification is hermetic and does not pollute

the host keyring.

Switches all `gpg` calls in the Node.js updater to `execFileSync` with argument arrays, eliminating shell-expansion attack surface during checksum verification.

Adds `package_spdx.go` tool that generates per-package SPDX JSON SBOMs, accepting `--package_name`, `--version`, `--purl`, `--sha256`, `--url`, `--id`, `--copyright`, `--generates`, `--homepage`, `--supplier`, and `--description` flags.

+1 more

WAS THIS USEFUL?  

Chainguard

Feedback ⓘ 1 RELEASE · 2026-08-31 NOTES ↗

Chainguard provides trusted open source artifacts for every layer of your modern software stack—containers, language libraries, and VM images.

> Chainguard launched Guardener, a GitHub App combining AI-driven Dockerfile migration, hardened SHA-pinned GitHub Actions, and signed-commit enforcement, while also shipping open-beta custom Rego pull policies, org-wide build pinning for Chainguard Libraries, a more robust Windows `chainctl` self-updater, and 21 new container images.

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

01 **Custom Rego pull policies in chainctl** NEW 95

Introduces custom container pull policies (open beta): organizations write Rego expressions submitted as YAML manifests, managed via `chainctl`, supporting parameters with per-binding values and an ENFORCE mode that blocks pulls. The new `chainctl policies custom validate` subcommand reports the exact line and column of a policy error before storing it, enabling CI validation of Rego policy manifests.

— Runnable subcommand plus mechanism, format and enforcement mode named

Week of 2026-08-31

02 **Build pinning for Chainguard Libraries** NEW 95

Enables build pinning for all organizations using Chainguard Libraries with upstream fallback across JavaScript, Python, and Java, recording the first-resolved copy of each package version to keep lockfiles and checksums stable. The new `chainctl libraries cache` command lets you review which package-version copies (upstream mirror or Chainguard rebuild) your organization holds, or opt out of build pinning.

Review which package-version copies your organization has pinned under Chainguard Libraries build pinning, or initiate an opt-out.

```
$ chainctl libraries cache
```

— Named command with example and clear opt-out mechanism Week of 2026-08-31

03 **Commit Verification policy enforcement** NEW 65

New Commit Verification capability enforces cryptographically signed commits against a per-repository policy, supporting keyless (Sigstore) signatures and static keys such as GPG.

— Mechanism and signing methods named, no setup steps given product docs

04 **.chainguard/ configuration model for Guardener** NEW

65

New `.chainguard/` configuration model enables Guardener GitHub App capabilities (Hardened Actions, Commit Verification) per repository via committed config files.

— Names exact config directory location, no schema detail product docs

05 **Windows support for chainctl self-update** IMPROVED

65

Extends `chainctl` self-update on Windows to replace the running executable correctly and verify release attestations before installing, without requiring administrator privileges.

— Describes fix and behavior, no exact invocation given Week of 2026-08-31

06 **Hardened Actions for GitHub workflows** NEW

60

New Hardened Actions capability recommends and migrates GitHub Actions to Chainguard's hardened, SHA-pinned equivalents, delivered via non-blocking PR review comments or scheduled/on-demand migration pull requests.

— Explains mechanism but no config or command to invoke product docs

07 **Grace-period container images removed** BREAKING

60

The container images `kube-conformance:1.32`, `kubernetes:1.32`, `nextcloud-server:31`, `rke2-runtime:1.32`, and `teleport:17` have passed their grace periods and are no longer available.

— Named removed images give clear migration prompt Week of 2026-08-31

THINNER COVERAGE BELOW

08 **AI-driven Dockerfile migration via chainctl agent** NEW

55

New `chainctl agent dockerfile` commands run AI-driven, iterative migration of Dockerfiles to Chainguard Containers locally.

— Names exact command but no flags or example shown product docs

09 **21 new container images added** NEW 55

Adds 21 new container images including `dynatrace-operator`, `netshoot`, `strimzi-drain-cleaner`, `redpanda-console`, `apache-polaris`, `nomad`, `kserve-pmmlserver`, `kserve-xgbserver`, `victorialogs-vlagent`, `mlrun-api-fips`, `mlrun-fips`, and `kube-logging-operator-fluentd-drain-watch`, with both standard and FIPS variants where noted.

— Lists many named images but no usage instructions Week of 2026-08-31

└──▷ **BREAKING ON UPGRADE**

! The container images `kube-conformance:1.32`, `kubernetes:1.32`, `nextcloud-server:31`, `rke2-runtime:1.32`, and `teleport:17` have passed their grace periods and are no longer available.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Google SecOps

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Meet your business challenges head on with AI and cloud computing services from Google, including security, data management, and hybrid & multi-cloud.

> Google SecOps shipped a new pipeline event-breaking processor for splitting multi-event log entries and added three Mandiant Frontline Threats rule packs to the Content Hub.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 **Unroll processor for pipeline event breaking** NEW 85

The Unroll processor splits log entries containing arrays or slices of events into individual log events before parsing and ingestion. When inputs are raw string payloads, a Transform processor (e.g. set(body, ParseJSON(body))) must be positioned immediately before the Unroll processor in the pipeline execution sequence.

Split multi-event JSON array log entries into discrete events in a pipeline — parse the raw body first, then unroll it before ingestion.

> 📍 In the Google SecOps pipeline editor, add a Transform processor with expression set(body, ParseJSON(body)), then add an Unroll processor immediately after it in the pipeline execution sequence.

— Names exact processors, ordering requirement, and runnable config example.

changelog-20260831-14fa9df4

THINNER COVERAGE BELOW

02 **Mandiant Frontline Threats rule packs for Linux, macOS, Google Cloud**

NEW 40

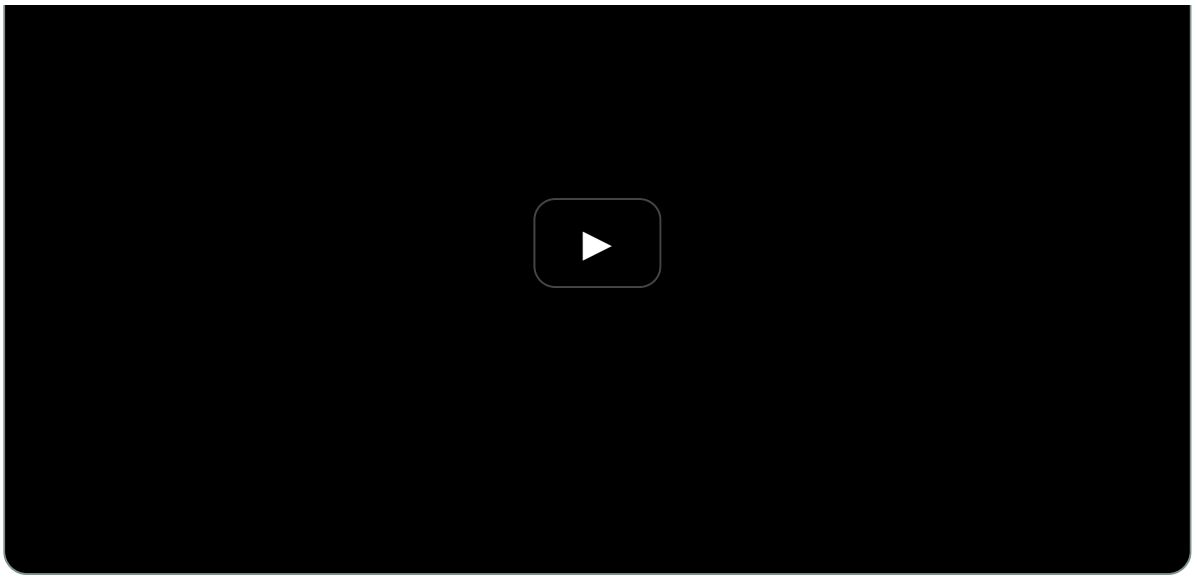
Three new rule packs were added to the Content Hub: Mandiant Frontline Threats for Linux, Mandiant Frontline Threats for MacOS, and Mandiant Frontline Threats for Google Cloud.

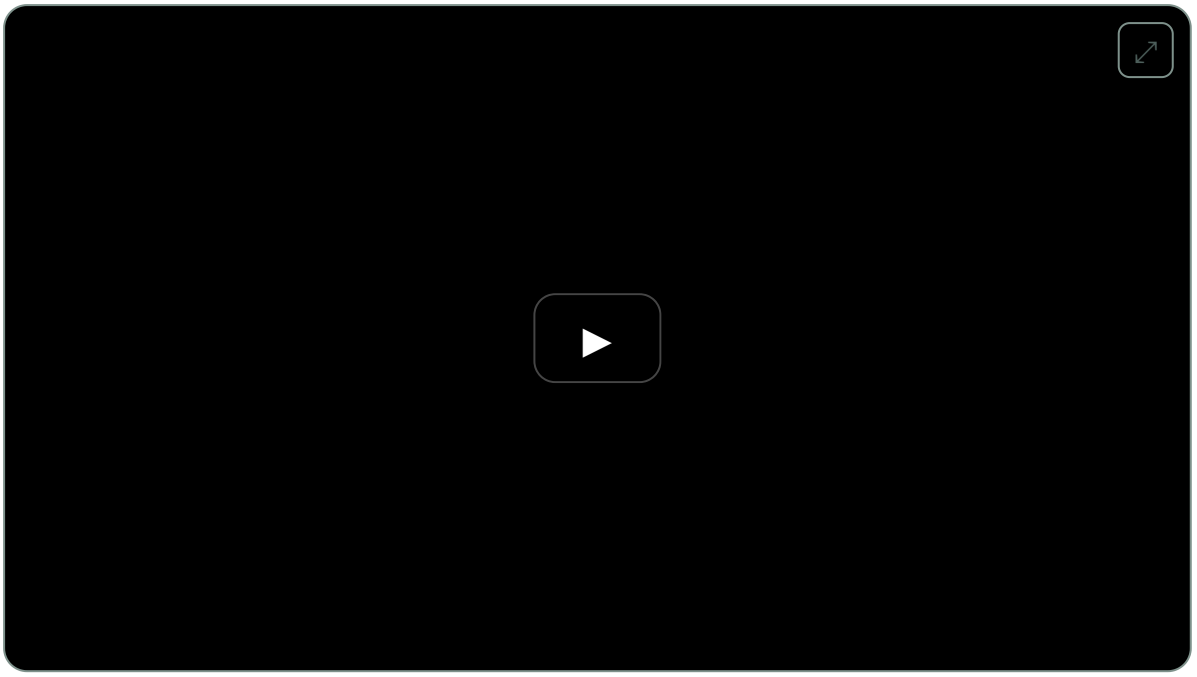
— Names all three packs but no detail on rule content or usage.

changelog-20260831-14fa9df4

└──▷ **ALSO FROM THESE RELEASES**







WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ 3 RELEASES · 2026-08-31 BLOG ↗

Sumo Logic is a cloud-based platform for real-time log management, monitoring, and security analytics across applications and infrastructure.

> Sumo Logic's biggest move this window is letting admins provision the SOC Analyst Agent directly in child org baselines with per-day investigation limits, alongside a batch of new integrations (Druva, four GA automation integrations, Terraform resources), source and API updates, and new auto-refresh guardrails for dashboards.

└──▷ **WHAT SHIPPED · 8 FEATURES** most completely described first ? what's the number?

01 **Two new Terraform provider resources** NEW 70

SumoLogic Terraform Provider v3.3.0 adds two new resources: `sumologic_async_aws_lambda_invocation` and `sumologic_s3_logging_lambda_enable`.

– Exact resource names given, but no usage or behavior detail.

Apps, Solutions, and Collection Integra…

02 **Auto-refresh guardrails for dashboards** NEW 70

Dashboards gain two auto-refresh guardrails: auto-refresh pauses when a tab is hidden for more than 10 minutes and resumes automatically on return, and a panel is dropped from auto-refresh after 5 consecutive query failures so it doesn't disrupt other healthy panels.

– Exact thresholds and behavior given, but no config toggle mentioned.

Auto-Refresh Guardrails for Dashboards …

03 **SOC Analyst Agent in child org baselines** NEW 60

Sumo Logic now lets admins provision the SOC Analyst Agent when creating or editing a child org baseline, alongside existing Logs, Metrics, Traces, and Cloud SIEM Enterprise baselines, with a configurable Number of investigations per day control and optional overage.

– Names the config control but gives no UI path or API detail.

SOC Analyst Agent Support in Child Org …

THINNER COVERAGE BELOW

04 **Miro app and source API version update** IMPROVED 55

Miro app and source updated to a new API version, removing dependency on the deprecated Audit Log API v2, with new out-of-the-box monitors added.

– Explains the migration away from a deprecated API, but not the new monitors.

Apps, Solutions, and Collection Integra…

05 **Automation Service Integrations reach GA** IMPROVED 45

Four Automation Service Integrations are now generally available: AWS Lambda, DFIR-IRIS, Google Cloud Functions, and LaunchDarkly.

— Lists integrations reaching GA, no setup detail given.

Apps, Solutions, and Collection Integra…

06 **New Druva Platform Events app** NEW 40

New Sumo Logic app for Druva - Platform Events enables monitoring of Druva via a webhook-based integration.

— Thin description, no config or endpoint detail.

Apps, Solutions, and Collection Integra…

07 **Akamai SIEM API parser path support** IMPROVED 40

Akamai SIEM API v1.2.26 adds parser path support.

— Names exact API version but no config or usage detail.

Apps, Solutions, and Collection Integra…

08 **PostgreSQL support for Google Cloud SQL source** NEW 35

Google Cloud SQL source adds support for PostgreSQL instances.

— Bare statement of new instance type support. Apps, Solutions, and Collection Integra…

WAS THIS USEFUL?  

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

3 RELEASES · 2026-08-14 → 2026-08-31 BLOG ↗

Trend Micro's unified security platform providing threat intelligence, endpoint protection, and attack surface visibility across cloud and on-premises environments.

> Trend Vision One shipped automatic disabling of inactive local user accounts with configurable thresholds and audit/notification support, alongside enforcement of Asset Visibility Management scopes within Endpoint Alerts.

└──▷ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Automatic disabling of inactive user accounts NEW

84

An Inactive user account auto-disable option in the User Accounts app Settings menu (top right of console) lets administrators automatically disable local accounts inactive for 30, 60, 90, or 180 days. Accounts disabled this way appear grayed out on the Select business page for multi-business users, and disabled users attempting to sign in see a clearer message directing them to contact their administrator. A new Status set by column in User Accounts shows whether an account's status was last changed by a user or the system, all auto-disable configuration changes and disable actions are recorded in the audit log, and an email notification is sent to a local user account when the system automatically disables it for inactivity.

Enable automatic disabling of local accounts that have not signed in for 90 days to enforce an idle-account policy.

> 📍 In the Trend Vision One console, go to Administration → User Accounts, open the Settings menu (top right), enable Inactive user account auto-disable, and select 90 days as the inactivity threshold.

Enable automatic disabling of inactive local user accounts to reduce your attack surface from dormant credentials.

> 📍 In the console, go to Administration → User Accounts, open Settings (top right), enable Inactive user account auto-disable, and select an inactivity threshold of 30, 60, 90, or 180 days.

— Names exact UI path, thresholds, and all related behaviors

Automatically disable inactive user acc… Automatically disable inactive user acc…

THINNER COVERAGE BELOW

02 **Asset Visibility Management scope in Endpoint Alerts** IMPROVED

50

Endpoint Alerts now enforces Asset Visibility Management scopes so users only see alerts and impacted endpoints within their assigned visibility scope, viewable under Endpoint Security → Endpoint Alerts.

Ensure your Asset Visibility Management scope is correctly configured so users see the right alerts in Endpoint Security → Endpoint Alerts.

> 📍 In the console, go to Asset Visibility Management and configure the visibility scope for each user or group, then verify results under Endpoint Security › Endpoint Alerts.

— Names the scope mechanism and UI path but lacks further detail

Endpoint Alerts applies Asset Visibilit…

WAS THIS USEFUL? 👍 👎

Corellium

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

Advice, answers, updates, and more from the Corellium Team

> Corellium launched support for the Arm RD-1AE automotive virtual hardware platform, bringing advanced hypervisor, safety-island, and trusted-firmware virtualization to its Arm Virtual Hardware offering.

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

01 **Arm RD-1AE automotive virtual hardware platform** NEW

80

Corellium added the Arm RD-1AE virtual hardware platform to the Arm Virtual Hardware platform powered by Corellium, virtualizing an advanced automotive reference design with Neoverse-V3AE application processors, Cortex-R82AE safety islands, and a Cortex-M55 Runtime Security Engine. The platform supports execution of Type 1 and Type 2 hypervisors, including the Xen hypervisor from Arm's reference software package, and includes SystemReady IR support, Arm Trusted Firmware, and trusted services. Pre-built RD-1AE software images covering the full set of Arm reference software deliverables are selectable at device creation time, and the platform boots to a Linux prompt in under 30 seconds on AWS Graviton 3 instances.

— Names processors, hypervisors, boot time; no exact command or endpoint given.

launch-20260831-aa561c03

WAS THIS USEFUL? 👍 👎

Sophos Mobile

Feedback

1 RELEASE · 2026-08-31

NOTES ↗

Sophos Mobile is an enterprise mobile device management and security platform that protects iOS and Android devices through threat prevention and policy enforcement.

> Sophos Mobile introduces Apple Declarative Device Management (DDM) policies for real-time desired-state configuration and clarifies domain filtering rules across Mobile Threat Defense and Chrome Security policies.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 **Web Filtering domain field guidance** IMPROVED 60

Improves guidance on the Web Filtering pages in Mobile Threat Defense policies for Android and iOS, and the Chrome Security policy for ChromeOS, clarifying that protocol prefixes such as `https://` or `chrome://` must not be included in the Allowed domains and Blocked domains fields.

— Names exact fields and prefixes, but is a documentation clarification

2026.35Released on August 31, 2026

THINNER COVERAGE BELOW

02 **Apple Declarative Device Management (DDM) policies** NEW 55

Adds support for Apple Declarative Device Management (DDM) policies, enabling desired-state configuration where devices self-maintain their configuration status and report changes immediately rather than waiting for periodic sync checks.

— Explains mechanism and behaviour but no config steps given

2026.35Released on August 31, 2026

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Beelzebub

Feedback ⓘ 1 RELEASE · 2026-08-31 NOTES ↗

A secure low code deception runtime framework, leveraging AI for System Virtualization.

> Beelzebub v3.9.1 adds method-aware HTTP routing to its honeypot service to make fingerprinting harder for scanners.

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

01 **Method-aware HTTP honeypot routing** IMPROVED 40

The HTTP honeypot service now routes and responds differently depending on the HTTP method used by a scanner, reducing the ability of scanners to fingerprint the honeypot.

— Explains behaviour change but no config/flag or endpoint named v3.9.1

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Censys

Feedback ⓘ 2 RELEASES · 2026-08-31 BLOG ↗

Censys is a search engine for discovering internet-connected devices, servers, and certificates to assess security exposure and vulnerabilities.

> Censys overhauled host reputation with an ML-based, SHAP-explainable scoring model and rolled out a broad ASM web-scanning upgrade adding screenshots, redirect tracking, JavaScript fingerprinting/risk detection, and a switch to ActiveDNS resolution.

└──▷ WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 ML-based host reputation scoring with SHAP explainability

72

BREAKING

Each host now receives a label (Benign , Suspicious , Malicious , Honeypot , or Inactive), a 0–100 score, a confidence level, and evidence for the top five signals that influenced the assessment, generated using a SHAP (SHapley Additive exPlanations) model. Some data fields from the legacy reputation score method changed as a result, so existing queries or integrations relying on those fields should be reviewed against Censys documentation before upgrading.

— Names label set and scoring mechanism but migration steps are vague

August 31, 2026 (2026-08-31)

02 ActiveDNS resolution for ASM web scanning IMPROVED

71

ASM web scanning switches to ActiveDNS resolution, increasing detected web entity counts by 0–10% while reducing false-positive host counts by 2–15%; host counts drop correspondingly because bare IPs not seen in the last 24 hours are no longer retained and only the most recent active resolution is shown.

— Gives concrete percentage impacts but resolution mechanics stay high-level

Guide to 2026 ASM Web Scanning Update (…

03 Forgejo detection in Platform and ASM NEW

65

Adds a new Forgejo software fingerprint to the Platform to detect self-hosted Forgejo Git forge instances, and an Exposed Forgejo Application risk fingerprint to ASM that flags exposed HTTP services running Forgejo due to the sensitive source code, credentials, and CI/CD configs they may contain. New risk types like this may be disabled by default, so workspaces should check risk type configuration in the ASM web console to enable them.

– Names both fingerprint and risk but no query examples August 31, 2026 (2026-08-31)

04 **Web entity instance limit to most recent resolution** BREAKING 64

Web entity instances are now limited to a maximum of one per web entity, representing only the most recent active IP resolution rather than all historical resolutions; as a result the `web_entity.instance_count` field will show significantly reduced values after migration, requiring review of any dashboards or alerting thresholds based on instance counts.

– Names the exact field affected and migration impact

Guide to 2026 ASM Web Scanning Update (…

05 **Automatic web screenshot capture in ASM** NEW 61

ASM now automatically captures web screenshots for every web entity in inventory, storing exact timestamps and a historical view over time so analysts can surface exposed login pages, admin panels, and potentially compromised infrastructure.

– Explains mechanism and use case, no direct API/UI path given

Guide to 2026 ASM Web Scanning Update (…

06 **JavaScript risk detection for EOL libraries** NEW 61

ASM adds JavaScript risk detection to identify and flag vulnerabilities in end-of-life JS libraries, enabling rapid response to JS-related incidents such as React2Shell. As with other new risk types, this may be disabled by default and require enabling via risk type configuration in the ASM web console.

– Names a concrete incident example but lacks scope of coverage

Guide to 2026 ASM Web Scanning Update (…

THINNER COVERAGE BELOW

07 **JavaScript library fingerprinting** NEW 46

ASM adds JavaScript library fingerprinting, using a large number of software fingerprints, to surface unknown or unauthorized JS dependencies across web assets.

– Describes capability but not how many fingerprints or where to view results

Guide to 2026 ASM Web Scanning Update (…

08 **Full redirect path tracking for web entities** NEW 42

ASM now tracks the full redirect path across web entities to help detect domain takeovers or incorrect redirect destinations.

– States purpose but no mechanism detail or navigation path

Guide to 2026 ASM Web Scanning Update (…

└──▷ **BREAKING ON UPGRADE**

! Some data fields from the legacy reputation score method changed as a result of the new ML-based reputation release — consult the Censys documentation to identify affected fields before upgrading queries or integrations that rely on them.

! The `web_entity.instance_count` field will have significantly reduced values after migration, as each web entity will have a maximum of one instance instead of multiple historical instances — any dashboards or alerting thresholds based on instance counts will need to be reviewed.

! Host counts in ASM inventory will drop by 2–15% as bare IPs that were not seen in the last 24 hours are no longer retained, and only the most recent active resolution is displayed.

WAS THIS USEFUL?  

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Aembit

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Aembit is a zero-trust identity security platform that manages and authenticates machine identities across cloud infrastructure.

> Aembit ships an in-process Edge SDK (TypeScript and Python) for proxy-free workload authentication, alongside new tracing fields in Agent Proxy events, installer property handling, a glibc-linked CLI build, and a consolidated Developer Guide.

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

01 **Edge SDK for TypeScript and Python** NEW 85

Introduces the Aembit Edge SDK (TypeScript and Python): in-process libraries that handle workload attestation, token lifecycle, and credential retrieval with no proxy required, suited for serverless functions, ephemeral CI containers, and restricted runtimes. It supports seven Trust Providers — AWS Metadata Service, AWS Role, Azure Instance Metadata Service, GCP Identity Token, GitHub, GitLab, and OIDC ID Token — with integration guides for AWS EC2, AWS Lambda, GCP Cloud Run, and Vercel OIDC.

— Detailed capability and named providers, but no runnable command shown.

changelog-20260831-f0b057c8

02 **Installer property handling for Windows MSI and Linux script**

IMPROVED 83

The Windows MSI installer now applies `AEMBIT_TENANT_GRPC_PING_INTERVAL_SECS` and `AEMBIT_TENANT_GRPC_PING_TIMEOUT_SECS` when passed as install-time properties. The Linux installer script now validates Agent Proxy arguments — including numeric environment variables such as `AEMBIT_HTTP_IDLE_TIMEOUT_SECS` — before installation proceeds.

Apply gRPC keep-alive settings at install time on Windows so the CLI maintains its tenant connection over idle periods.

```
$ msiexec /i AembitCLI.msi
AEMBIT_TENANT_GRPC_PING_INTERVAL_SECS=30
AEMBIT_TENANT_GRPC_PING_TIMEOUT_SECS=10
```

— Named env vars paired with a runnable install command. changelog-20260831-f0b057c8

03 **New fields in Agent Proxy workload and response events**

IMPROVED

70

Adds `authorizationChain` to Agent Proxy workload events — an ordered list of context IDs for the authorization events behind a request, enabling full trace-back to the directive and credential retrievals. Also adds `severity`, `outcome.result`, and `outcome.reason` to response events when Aembit answers in place of the Server Workload, so the outcome is no longer cleared between request and response.

— Named event fields with mechanism, but no usage example given.

changelog-20260831-f0b057c8

04 **glibc-linked Linux amd64 CLI build** NEW

70

Aembit CLI 1.34.5772 adds `aembit_agent_cli_linux_amd64_glibc_<version>.tar.gz` as a second Linux amd64 archive, dynamically linking against glibc 2.28 or newer, for organizations whose policy requires glibc or to work around DNS resolution issues caused by musl's stricter handling of resolver responses.

— Named artifact and glibc version, no usage example. changelog-20260831-f0b057c8

THINNER COVERAGE BELOW

05 **Unified Developer Guide documentation** IMPROVED

55

Consolidates API Guide, CLI Guide, and Edge SDK documentation into a single Developer Guide under `/dev-guide/`, with a side-by-side credential-path comparison of Agent Proxy, Edge SDK, Aembit CLI, and Edge API.

— Names the doc path but describes only reorganization. changelog-20260831-f0b057c8

WAS THIS USEFUL? 👍 👎

Didit

Feedback

ⓘ

2 RELEASES · SEEN 2026-08-31

NOTES ↗

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit shipped a Marketplace for swapping identity-verification engines per workflow feature with automatic fallback, a full biometric-template retention system spanning webhook config, per-call overrides, a new Biometric Templates API and privacy-erasure controls, plus transaction-monitoring rule management via REST and AI-assisted rule migration through MCP. Note a breaking change to the session-delete response format.

└──▷ **WHAT SHIPPED · 8 FEATURES** most completely described first ? what's the number?

01 **Third-party verification engine marketplace** NEW 95

Adds `provider_key` (string) and `fallback_to_native` (boolean, default `true`) fields to the feature `config` object in the Management API for programmatic engine selection and fallback control. A new Marketplace UI lets you connect third-party identity verification providers and choose, per workflow feature node, which engine runs the check via a Verification engine control on the node's first configuration tab, with status labels Active, Switch, and Connect and inline provider setup without leaving the workflow builder. A 'Fall back to Didit if my provider fails' toggle (default on) automatically runs Didit's own engine if the selected third-party provider is unavailable.

Override the verification engine for a specific workflow feature and disable silent fallback, so failures surface explicitly rather than silently running on Didit.

json

```
{
  "config": {
    "provider_key": "your-provider-key",
    "fallback_to_native": false
  }
}
```

Route a workflow feature to a connected third-party provider while keeping the default automatic fallback to Didit if the provider is unavailable.

```json
{
  "config": {
    "provider_key": "your-provider-key",
    "fallback_to_native": true
  }
}
```

Connect a new provider and activate it for a feature without leaving the workflow builder.

> 📍 1. Open the workflow builder and select the feature node you want to change. 2. On the first configuration tab, open the `Verification engine` control. 3. Find the provider listed as `Connect` and select it — the setup dialog opens inline. 4. Complete the connection, then select the provider again (now shown as `Switch`) to make it the active engine for this feature.

Switch a verification step to a third-party provider so it runs through your connected Marketplace integration, with Didit as the automatic fallback.

> 📍 In the workflow editor, open a feature node › first configuration tab › set `Verification engine` to your connected third-party provider from the Marketplace.

— Names exact API fields, UI control and step-by-step activation path `product docs`

02 **Biometric-template retention configuration** NEW 95

Adds `face_retention_policy` and `face_retention_days` fields on `PATCH /v3/webhook/` to configure biometric-template retention at the application level, keeping image-free face embeddings alive for duplicate detection, Face Search, and biometric authentication after session data is purged. A `retain_face_embeddings` parameter on `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` overrides that policy per call, and a `Retain biometric template` toggle with a required finite retention period is available in App Settings → Data.

Enable biometric-template retention org-wide so Face Search keeps working after sessions are deleted, with a 90-day expiry.

```
$ curl -X PATCH https://api.didit.me/v3/webhook/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"face_retention_policy": "retain", "face_retention_days": 90}'
```

Override the app policy and suppress template retention for a single sensitive session deletion.

```
$ curl -X DELETE
https://api.didit.me/v3/session/{sessionId}/delete/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"retain_face_embeddings": false}'
```

— Endpoint fields plus runnable curl examples for both levels `snapshot-20260831`

---

03 **Biometric Templates management API** NEW 85

New Biometric Templates API — `GET /v3/biometric-templates/`, `GET /v3/biometric-templates/count/`, `GET /v3/biometric-templates/{template_uuid}/`, `DELETE /v3/biometric-templates/{template_uuid}/`, and `POST /v3/biometric-templates/delete/` — lists, counts, inspects, and purges retained templates, with every retain and purge action audited.

— Five endpoints named, but no example call given `snapshot-20260831`

---

04 **Privacy erasure and cascading purge on deletion** NEW 83

Adds `deletion_instruction: "privacy_erasure"` (with `instruction_id`) on session deletion to purge all retained biometric templates for a person regardless of application policy. `POST /v3/users/delete/` now purges a user's retained biometric templates before deleting the user, returning `503` (retryable, nothing deleted) if a template cannot be purged.

Purge all retained biometric templates for a person on a GDPR erasure request, regardless of application retention policy.

```
$ curl -X POST https://api.didit.me/v3/sessions/delete/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"deletion_instruction": "privacy_erasure",
"instruction_id": "<your-instruction-id>"}'
```

— Named fields and endpoints with a runnable erasure example `snapshot-20260831`

---

05 **Transaction-monitoring rules API and library** NEW 80

New `/v3/transactions/rules/` Management API endpoints let you create, update, delete, and backtest transaction-monitoring rules programmatically, with matching

`library` and `install` endpoints for browsing, installing, and removing preset rules from the rule library.

— Endpoint surface named, no example request shown `changelog-20260831-93a5934d`

---

06 **Session-delete response now returns 200 with body** BREAKING

78

`DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body (`face_retention_outcome` , `biometric_template_uuid` ; per-session `results[]` for batch) instead of `204 No Content` — clients that assert on a `204` response will break and must be updated.

— Exact before/after response contract named, no code migration sample `snapshot-20260831`

---

07 **AI-assisted rule migration via MCP** NEW

65

New AI-assisted rule migration workflow lets the Didit assistant, or any MCP-connected agent, map a rule set exported from another provider (CSV or Excel) onto Didit's rule schema, backtest it against your transaction history, and create the rules in test mode for review before activation.

— Mechanism described in prose but no command or endpoint given

`changelog-20260831-93a5934d`

THINNER COVERAGE BELOW

---

08 **Face Search retained-template match metadata** IMPROVED

50

Face Search match results now include `source: "retained_template"` , `vendor_user_id` , and `biometric_template_id` fields for hits on retained templates, and `vendor_user_id` on session and imported matches.

— Field names given but no example response or usage steps `snapshot-20260831`

└──▷ **BREAKING ON UPGRADE**

! `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body (`face_retention_outcome` , `biometric_template_uuid` ; per-session `results[]` for batch) instead of `204 No Content` — clients that assert on a `204` response will break and must be updated.

WAS THIS USEFUL? 👍 👎

Material Security

Feedback

1 RELEASE · SEEN 2026-08-31

BLOG ↗

Material Security is an email security platform that uses AI to detect and prevent advanced threats like phishing, ransomware, and data exfiltration.

> Material Security overhauled Google Drive detection and response with more granular sensitive-content detections, natural-language detection creation, configurable remediation flows, in-context manual fixes, and a new per-file audit log.

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

01 **Granular sensitive-content detection categories** IMPROVED 57

Previously combined sensitive-content detections are now split into distinct categories — passwords, payroll information, and social security numbers each produce their own individual detection — giving analysts a clearer, more actionable risk view.

— Names three new categories but no config surface given launch-20260831-a87508da

02 **Natural language detection creation** NEW 55

Natural language search is extended to detection creation, letting analysts describe what they want to detect in plain English and activate it as a live detection without writing MQL or building a structured query.

— Clear capability but no example query or UI path given launch-20260831-a87508da

03 **Configurable remediation flows for overpermissioned files** NEW 55

Security teams can now set configurable remediation flow options for overpermissioned files, including the order of notifications, grace periods, and remediation actions, to match organizational sharing policies.

— Describes mechanism but no config key or exact settings named launch-20260831-a87508da

04 **Per-file audit log** NEW 47

A new per-file audit log shows every security action taken on a file and when, with forthcoming support for analyst comments on the audit timeline and clearer error messages when an automated remediation step fails.

— Describes feature and roadmap but no access path or format launch-20260831-a87508da

05 **Manual remediation from file pages** IMPROVED 45

Manual remediation actions are now available directly on file issue pages and file detail pages, removing the need to navigate to the detection page to fix over-sharing.

— Names UI location but not exact navigation steps `launch-20260831-a87508da`

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Secure users, AI agents, and more with Auth0, an easy-to-implement, scalable, and adaptable authentication and authorization platform.

> Auth0 brought Custom Token Exchange to General Availability with full Dashboard management, added shared transaction context between exchange and post-login Actions, and enabled session delegation for AI and support agents; WebAuthn biometrics can now serve as a first authentication factor.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Shared transaction context in Custom Token Exchange** NEW 45

Adds shared transaction context between Custom Token Exchange Actions and Post-Login Actions via transaction metadata, enabling token customization based on subject or actor token information.

— Names the transaction metadata mechanism but no config steps shown

changelog-20260831-fac36582

02 **Session Delegation via Custom Token Exchange** NEW 40

Enables Session Delegation through Custom Token Exchange, allowing a service, AI agent, or support agent to establish a web session on behalf of a user.

— Describes the use case but not the API or config to enable it

changelog-20260831-fac36582

03 **WebAuthn biometrics as a first factor** NEW 20

WebAuthn biometrics can now be used as a first factor for authentication, alongside the Custom Token Exchange GA release.

— Only mentioned in passing, no mechanism or setup detail given

changelog-20260831-fac36582

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

Learn what password cracking is, how it works, and the best ways to protect against it.

> JumpCloud launched Access Risk Detection, a new identity threat capability that builds per-user behavioral baselines and surfaces risky logins through weighted scoring and a review dashboard.

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

01 Access Risk Detection behavioral anomaly scoring NEW

77

A new Access Risk Detection capability builds a unique behavioral profile per user from 60 days of continuous activity or the last 50 successful logins, then correlates concurrent signals (e.g., unfamiliar browser paired with impossible travel) into a single weighted risk score. Events are mapped into three severity tiers — Low, Medium, and Critical — and privileged account weighting applies a higher score to deviations on admin profiles than on standard user accounts. Automated step-up MFA verification loops trigger additional authentication when anomalous login signals are detected, reducing false positives. A dedicated Needs Review dashboard filters out auto-resolved events, surfaces granular telemetry, supports custom admin filters, and allows compliance notes to be attached directly to an incident.

— Rich mechanism and named dashboard, but no API or config key given

launch-20260831-7a6ed76d

WAS THIS USEFUL? 👍 👎

CrowdStrike Falcon Identity Protection

Feedback

ⓘ 1 RELEASE · 2026-08-31 PRESS ↗

CrowdStrike Falcon Identity Protection detects and prevents identity-based attacks through behavioral analytics and threat intelligence across enterprise environments.

> CrowdStrike Falcon Identity Protection added a real-time biometric verification integration with CLEAR, correlating identity-proofing results with other Falcon signals to drive access decisions.

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

01 **CLEAR biometric identity verification integration** NEW

53

Falcon now integrates with CLEAR to trigger real-time biometric person-verification when identity risk is detected, correlating CLEAR identity verification results alongside other Falcon platform signals to determine whether activity should be allowed, investigated, or blocked.

— Explains mechanism and decision outcomes but no config or API detail.

2026-08-31-ce30cf5f

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback

ⓘ

1 RELEASE · 2026-08-31

PRESS ↗

SailPoint is an identity governance platform that manages user access, permissions, and compliance across enterprise applications and systems.

> SailPoint launched SecOps Identity Intelligence integration with CrowdStrike Falcon Next-Gen SIEM, surfacing identity governance and risk data inside SOC investigation workflows.

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

01 **SecOps Identity Intelligence integration with Falcon Next-Gen SIEM**

NEW

48

SailPoint identity governance, access, and risk data for human, machine, and AI agent identities is brought into CrowdStrike Falcon Next-Gen SIEM for use during security investigations, allowing correlation with endpoint, cloud, threat intelligence, and third-party telemetry already present in the SIEM.

— Describes purpose and data types but no config, endpoint, or setup steps.

launch-20260831-0d4a3bff

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ 1 RELEASE · 2026-08-31

CHANGELOG ↗

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across infrastructure.

> CertCentral now lets administrators configure the default Key Usage pre-selection on the X9 PKI for TLS certificate request form.

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

01 Configurable default Key Usage for X9 PKI for TLS NEW

62

A new Default Key Usage setting on the X9 PKI for TLS Product Settings page lets CertCentral Administrators choose whether Digital signature or Digital signature and key encipherment is pre-selected on the X9 PKI for TLS certificate request form.

— Names exact setting, page, and two options; no broader mechanism given.

changelog-20260831-ca5f15c3

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

Give feedback to the DNSFilter team so we can make more informed product decisions. Powered by Canny.

> DNSFilter made DNS PreCheck v2 the default anti-hijacking mode for Windows Roaming Clients, added a query log exclusion filter, moved Relay Manager's config file to a safer default path, and rolled out several dashboard usability upgrades.

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

01 **DNS PreCheck v2 default for Windows Roaming Clients** IMPROVED

67

DNS PreCheck v2 is now the default filtering mode for all Windows Roaming Clients on v3.6 and later, detecting and overriding network-level DNS hijacking such as hotel Wi-Fi, captive portals, and intercepting ISPs; devices with it active show `HTTP/2.0` protocol in the DNS query log.

— Explains mechanism and behaviour but only a UI-visible signal, not a config path.

changelog-20260831-e285616c

02 **'Does Not Include' operator in Query Log** NEW

63

Adds a Does Not Include operator to the Categories filter in the Query Log, letting admins exclude specific categories from view during investigations.

Narrow a query log investigation to exclude noisy categories (e.g., hide all News traffic) so only relevant results remain.

> 📍 In the DNSFilter portal, go to Insights › Query Log, open the Categories filter, select the Does Not Include operator, and choose the categories to suppress (e.g., News).

— Comes with an exact navigation path and worked example. changelog-20260831-e285616c

THINNER COVERAGE BELOW

03 **Relay Manager config path moved to system location** IMPROVED

53

Relay Manager v0.2.2 changes the default config file installation path to a shared system location on Windows, preventing configuration loss when user profiles are removed.

— Names the mechanism and benefit but not the exact path. changelog-20260831-e285616c

04 **Customizable Dashboard usability upgrades** IMPROVED

47

In Manage Dashboards, dashboards now display Private or Shared visibility status, support updates from the quick actions menu, and are clickable from the grid; the Total Requests widget now opens Insights with Allowed Requests pre-toggled on.

— Lists several UI tweaks grouped under one area without named config surfaces.

changelog-20260831-e285616c

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale launched `tailcat`, a new netcat-like CLI tool and open-source Go package that tunnels bidirectional bytes over its WireGuard-based magicsock data plane without requiring a Tailscale account.

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

01 tailcat: netcat-like tool over Tailscale's WireGuard data plane

NEW 70

Tailscale released `tailcat`, a new CLI tool and open-source Go package that lets two machines exchange bidirectional bytes over Tailscale's magicsock (WireGuard + NAT traversal + DERP relay) with no Tailscale account required, no IP addresses visible to users, and no OS-level networking changes (no TUN devices, no routing table modifications).

– Explains mechanism and scope but gives no runnable command launch-20260831-eca3abff

02 Client mode with transparent SOCKS proxy for child processes

NEW 70

Client mode supports running a child process (e.g. `curl`) with an ephemeral local SOCKS proxy exposed via an environment variable, letting `tailcat -oblivious` programs use the encrypted tunnel transparently.

– Names mechanism and example program but not the exact env var launch-20260831-eca3abff

03 Custom DERP server support with automatic fallback NEW

70

`tailcat` supports specifying a custom DERP server, such as a self-hosted `cmd/derper` instance, for full independence from Tailscale infrastructure, and falls back automatically to Tailscale-hosted, rate-limited DERP relays when direct UDP NAT traversal fails.

– Names self-host mechanism and fallback behavior, no exact flag launch-20260831-eca3abff

THINNER COVERAGE BELOW

04 Key locking and keypair reuse in tailcat NEW

40

tailcat supports optional key locking so the server can restrict which client public keys are permitted to connect, and keypairs can be ephemeral or named and reused across sessions.

— Thin one-line descriptions of two key-related options `launch-20260831-eca3abff`

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Varonis

Feedback ⓘ 1 RELEASE · 2026-08-31 CHANGELOG ↗

The leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis shipped AWS Organization-scoped scan targeting, a unified Atlassian (Jira+Confluence) integration, scheduled scan-pause windows, a new Data Classes classification building block, and faster Snowflake sharing and masking-policy visibility.

└──▷ WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 **CSV export of connected service accounts** NEW 65

Adds CSV export of filtered connected service accounts directly from the Connected Services page, with support for advanced filter criteria (e.g. by service type or last-seen date) before exporting.

Export a filtered list of connected service accounts to CSV for auditing large environments from the Connected Services page.

> 📍 In the Varonis console, go to Connected Services › Connected Service Accounts, apply your advanced filter criteria (e.g. by service type or last-seen date), then select Export › CSV to download the filtered results.

— Names exact UI path and export format but no limits given changelog-20260831-e1a498a6

02 **Scheduled scan pauses by day and time zone** NEW 65

Adds scheduling of scan pauses by day of the week and time zone, configured under Settings > Scan Management > Scan Pauses, to align maintenance windows across distributed environments.

Align scan pauses to business hours and maintenance windows across time zones by configuring per-day pause schedules.

> 📍 In the Varonis console, go to Settings › Scan Management › Scan Pauses, select a day of the week, set the pause window start and end times, and choose the applicable time zone. Repeat for each day as needed.

— Exact navigation and scheduling parameters given changelog-20260831-e1a498a6

03 **AWS Organization ID as scan scope target** NEW 65

Adds AWS Organization ID as a scan scope rule target, allowing entire AWS Organizations to be included or excluded without maintaining per-account lists, configured under

Settings > Scan Scopes.

Scope a classification scan to an entire AWS Organization to avoid maintaining per-account inclusion lists.

> 📍 In the Varonis console, go to Settings › Scan Scopes, create or edit a scope rule, and enter your AWS Organization ID as the target to include or exclude that Organization from classification scanning.

— Exact config location and mechanism given `changelog-20260831-e1a498a6`

THINNER COVERAGE BELOW

04 **Data Classes for classification policies** NEW 55

Adds Data Classes as a building block for classification policies, automatically mapping related patterns and locale-specific criteria to reduce complexity and false positives.

— Mechanism described but no config path or example `changelog-20260831-e1a498a6`

05 **Snowflake sharing visibility and faster change sync** IMPROVED 55

Adds visibility into why a Snowflake resource is labeled Shared Internally or Shared Externally, showing originating shares and associated accounts, and reduces Snowflake masking policy and tag change visibility latency to within 24 hours instead of waiting for the next scheduled sync.

— Concrete latency figure but no example or UI path `changelog-20260831-e1a498a6`

06 **Unified Atlassian integration for Jira and Confluence** NEW 50

Adds a unified Atlassian integration that onboards Jira and Confluence together with shared identity, eliminating duplicate users.

— Names both apps but no setup steps given `changelog-20260831-e1a498a6`

07 **Data coverage reporting across scan scopes** NEW 40

Adds data coverage reporting across scan scopes, including historical trends and breakdown by cloud provider and service.

— Describes scope of report but no access path `changelog-20260831-e1a498a6`

08 **Pending automation action email notifications** NEW 30

Adds email notifications when automation actions enter a pending state requiring review in Action Center.

— Bare description with no configuration detail `changelog-20260831-e1a498a6`

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

Jazz Security

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

Jazz Security is an AI-native data loss prevention platform that investigates risky data movement and insider activity.

> Jazz launched a native CrowdStrike Foundry app that routes Falcon sensor data-protection signals through its Melody agentic investigator, auto-closes findings that need no human attention, and can be extended with an optional endpoint agent for deeper visibility.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Jazz app in CrowdStrike Foundry catalog** NEW 62

Jazz's Melody agentic investigator is now installable as a native app in the CrowdStrike Foundry app catalog, routing Falcon sensor data-protection signals through Melody's investigation pipeline and returning pre-investigated findings directly into the Falcon console. Investigated findings cover GenAI data risk, departing-employee data movement, and shadow IT sprawl, derived from Falcon telemetry without requiring custom rules or regex.

– Names catalog and console integration but no setup steps launch-20260831-6461c70a

02 **Optional endpoint agent for extended coverage** NEW 60

An optional Jazz endpoint agent extends Melody's coverage beyond the Falcon sensor to prompt-level GenAI visibility, desktop applications, shadow IT, and personal cloud accounts, with under 1% CPU overhead.

– Concrete scope and overhead figure but no install instructions launch-20260831-6461c70a

THINNER COVERAGE BELOW

03 **Automatic closure of low-risk findings** NEW 40

Melody automatically closes findings that require no human attention, with every closure documented and evidence attached, directly within the Falcon analyst workflow.

– Clear behaviour described but no config or trigger detail launch-20260831-6461c70a

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

RegScale is a compliance automation and GRC platform that helps organizations manage regulatory requirements, assessments, and security governance.

> RegScale's CLI adds milestone-level verification to JCAM POA&M dry-run imports, letting practitioners validate mappings before committing data.

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

01 **Milestone mapping in JCAM dry-run imports** IMPROVED

70

The `jcam dry_run` command now follows POA&Ms through to their milestones, letting practitioners verify milestone mapping before a live import rather than after. By default it samples a subset of POA&Ms per system during the dry run, proving the mapping without enumerating every record.

— Names exact command and sampling behavior but no sample size or config given.

CLI 6.43.0 (2026-08-31)

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 2 RELEASES · 2026-08-31 CHANGELOG ↗

Cloudflare is a web performance and security platform providing DDoS protection, CDN, DNS, and developer tools.

> Cloudflare shipped Content Signals crawl enforcement for Browser Rendering, broadened MCP protocol support, raised Durable Object concurrency limits, added new Workers AI models including multimodal GLM-5.3-flash, and expanded Access service token lifecycle controls, while confirming the Azure Sentinel Logpush connector's retirement.

└──▷ WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 Content Signals crawl enforcement via `contentUse` NEW

95

The `POST /client/v4/accounts/{account_id}/browser-rendering/crawl` endpoint now accepts a `contentUse` parameter (`reference` or `full`, default `full`) to enforce Content Signals `use` directives declared in a target site's `robots.txt`; if the site's declared level is more restrictive than the requested `contentUse`, the request is rejected with a `400`.

Crawl a page for reference-only use, so the request is rejected if the site's robots.txt restricts content to below `reference` level.

```
$ curl -X POST
'https://api.cloudflare.com/client/v4/accounts/{account_id}/brows
er-rendering/crawl' \
  -H 'Authorization: Bearer <apiToken>' \
  -H 'Content-Type: application/json' \
  -d '{"url": "https://example.com", "contentUse": "reference",
"formats": ["markdown"]}'
```

— Exact endpoint, parameter values and error code given. changelog-20260831-caaa810f

02 `retryAlarm` option for Durable Object aborts NEW

85

ctx.abort() in Durable Objects now accepts a `{ retryAlarm: false }` option, preventing alarm retries after the object resets — useful for one-shot cleanup alarms.

Prevent a one-shot cleanup alarm from retrying after it aborts the Durable Object.

```javascript
this.ctx.abort("Cleanup complete", { retryAlarm: false });
```

— Exact API option with a runnable code example. changelog-20260831-a37144d7

---

03 **New Workers AI and AI Search models** NEW 80

Workers AI adds `@cf/zai-org/glm-5.3-flash` (first multimodal GLM-5 model, 320B total / 18B active MoE) and `@cf/zai-org/glm-5.3` (agentic coding model, priced at $1.40/M input tokens, $0.26/M cached input tokens, $4.40/M output tokens). AI Search gains six additional Workers AI text-generation models including `@cf/openai/gpt-oss-120b` (128,000-token context), `@cf/openai/gpt-oss-20b` (128,000-token context), and `@cf/moonshotai/kimi-k2.7-code` (262,144-token context), with context windows up to 1,048,576 tokens across the set; no additional provider key is required.

— Names every model and pricing/context but no direct call shown.

changelog-20260831-caaa810f changelog-20260831-a37144d7

---

04 **Grace period and disable for Access service tokens** NEW 80

Cloudflare Access administrators can set a grace period (1 hour to 30 days via dashboard, or a custom RFC 3339 expiration via the API) when rotating a service token secret, keeping both secrets valid during the transition, and can temporarily disable a service token without deleting it — a disabled token cannot authenticate, any active rotation grace period stops, and the configuration is preserved for re-enabling.

Temporarily disable a suspected-compromised Access service token without deleting it, preserving config for later re-enabling.

> 📍 Settings › Access › Service Tokens › select token › Toggle `Enabled` to off ('Turn a service token on or off')

— Exact time ranges and UI steps for disabling a token. changelog-20260831-caaa810f changelog-20260831-a37144d7

05 **Azure Sentinel Logpush connector deprecated** DEPRECATED

75

The Azure Functions-based Microsoft Sentinel Logpush connector stops being maintained after 2026-09-14, when Microsoft retires the Azure Monitor HTTP Data Collector API; affected Enterprise customers must migrate to the Cloudflare for Microsoft Sentinel Codeless Connector Framework (CCF) connector before that date.

— Exact deadline and migration path named. changelog-20260831-caaa810f changelog-20260831-a37144d7

06 **App-scoped API tokens for Flagship** NEW

70

Custom API tokens can now be scoped to specific Flagship apps via the Specified Flagship apps option in the resource dropdown, granting only Evaluate, Read, or Write Flagship App permissions per selected app.

— Names exact UI option and permission levels. changelog-20260831-caaa810f changelog-20260831-a37144d7

07 **MCP server supports 2026-07-28 protocol** IMPROVED

65

The MCP server portal's `/mcp` endpoint now accepts stateless MCP `2026-07-28` specification requests alongside earlier 2025 Streamable HTTP clients, with independent protocol negotiation for client and upstream connections; SSE connections continue to use the legacy protocol.

— Names endpoint and protocol versions but gives no runnable example. changelog-20260831-caaa810f changelog-20260831-a37144d7

THINNER COVERAGE BELOW

08 **Radar Researcher links and URL Scanner AI explanations** NEW

45

Radar Researcher responses can now link to relevant Radar pages, reports, and Cloudflare Blog posts, and the Explain with AI button on URL Scanner reports enables AI-driven explanation of scan findings with follow-up Q&A.

— Describes the feature but no config or endpoint given. changelog-20260831-caaa810f

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The Azure Functions-based Microsoft Sentinel Logpush connector stops being maintained after 2026-09-14 when Microsoft retires the Azure Monitor HTTP Data Collector API; affected Enterprise customers must migrate to the Cloudflare for Microsoft Sentinel Codeless Connector Framework (CCF) connector before that date.

! The Azure Functions-based Microsoft Sentinel Logpush connector will stop working after 2026-09-14 (when Microsoft deprecates the Azure Monitor HTTP Data Collector API); Enterprise customers must migrate to the Cloudflare for Microsoft Sentinel Codeless Connector Framework (CCF) connector before that date.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 31, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ **Subscribe** Feedback on this issue