THE CYBER TOOLCHAIN

NO. 053

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 22, 2026 · EVERY DAY

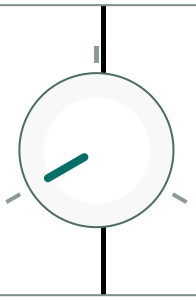
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 18 tools

APPLICATION SECURITY

vim ? help

◆ SAST / SCA / Code Security

Retire.js

Feedback ⓘ 3 RELEASES · 2026-08-21 → 2026-08-22 NOTES ↗

scanner detecting the use of JavaScript libraries with known vulnerabilities. Can also generate an SBOM of the libraries it finds.

> Retire.js expanded its SBOM output across three releases, adding CycloneDX 1.7 support with SPDX licensing and detection evidence, plus new metadata fields for certificate-check bypasses and vulnerability repository provenance.

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

01 **CycloneDX 1.7 SBOM output with SPDX licensing and evidence**

NEW 95

Adds `--outputformat cyclonedxJSON1_7` and `--outputformat cyclonedxJSON1_7_VEX` flags for CycloneDX 1.7 SBOM output. CycloneDX 1.6 and 1.7 reports now include `$schema`, `metadata.tools` (object form from spec 1.5), `metadata.component` (scanned path as BOM subject), `dependencies` (scan target linked to all detected components), `evidence.identity` (version-detection technique and confidence, e.g. `filecontent`, `filename`, `hash`, `ast`), and `vulnerabilities[].detail` (VEX variants only, when the repository supplies `details`).

Generate a CycloneDX 1.7 SBOM with full per-component SPDX license expressions and detection evidence for a Node.js project.

```
$ retire --outputformat cyclonedxJSON1_7
```

Generate a CycloneDX 1.7 VEX report (includes vulnerability detail fields) and exit 0 even if vulnerabilities are found, for use in CI pipelines that handle VEX downstream.

```
$ retire --outputformat cyclonedxJSON1_7_VEX --exitwith 0
```

— Names exact flags, fields and runnable commands 5.7.0

02 **Vulnerability repository metadata in reports** NEW 76

JSON reports gain a `vulnerabilityRepositories` field listing the vulnerability repositories used during the scan. All CycloneDX report formats gain `retirejs:vulnerability-repository` entries under `metadata.properties`, and CycloneDX 1.6 VEX output now identifies Retire.js as the source of each vulnerability and rating, including the repository URL when a single repository is used.

— Names fields and behaviour but lacks a runnable example 5.5.0

03 **Certificate bypass flagged in reports** NEW 68

When `--insecure` is set, JSON reports now include an `ignoreRepositoryCertificateErrors` field, and all CycloneDX reports include a `retirejs:ignore-repository-certificate-errors` entry under `metadata.properties`, making repository certificate-check bypasses machine-readable for auditors.

— Names exact fields but no usage example given 5.6.0

WAS THIS USEFUL? 👍 👎

GitHub

Feedback

3 RELEASES · 2026-08-21

BLOG ↗

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

> GitHub's window centered on bringing Copilot's agentic coding to chat platforms — Slack and Microsoft Teams both gained collaborative, async agent sessions that can open pull requests — alongside improvements to blocked-user moderation tools.

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

01 GitHub Copilot agent sessions in Microsoft Teams NEW

88

Mentioning `@GitHub` in Microsoft Teams channels, threads, or direct messages starts a Copilot cloud agent session visible to all conversation participants; participants with write access to the target repository can trigger code changes directly from the conversation. Work runs asynchronously in a secure cloud sandbox with progress visible in the thread, and artifacts can be picked up in a terminal, the GitHub Copilot app, or an IDE. Repository admins can require an additional approval before any pull request attributed to the Teams Copilot identity can merge; AI credit usage is governed by usage-based billing budgets (cloud sandbox usage billed separately, controllable at product-level or SKU-level). Public channels can be configured with a default repository (DMs cannot), and `@GitHub help` lists available commands.

— Names exact mention command, help command, and billing controls.

Shared agentic work with GitHub Copilot…

02 GitHub Copilot agent sessions in Slack NEW

82

Mentioning `@GitHub` in any Slack direct message, channel, or thread starts a Copilot agent session that can triage issues, implement changes in a secure cloud sandbox, and open pull requests; Copilot creates a dedicated Slack Code channel per task where teammates can inspect diffs and HTML artifact previews. Sessions are multiplayer — any team member can join from the original thread, add context, redirect, or stop it — and run asynchronously, with handoff to a terminal, the GitHub Copilot app, or an IDE via the resulting pull request. PRs and issues are attributed to the Copilot app identity and bounded by existing GitHub permissions; admins can require additional human approval before agent-authored PRs merge. Available in public preview to GitHub Copilot Business and Enterprise plans, with usage counted against existing Copilot entitlements and cloud agent budgets.

— Explains full mechanism and plans but no API/flag detail.

The new GitHub Copilot experience in Sl…

03 Search, filtering, and notes for blocked users IMPROVED

67

Blocked-user lists (Settings → Moderation → Blocked users, for personal and org accounts) now support searching by username, full name, or email; filtering by block reason (Spam, Misconduct, or Harassment); and sort and pagination for long lists. Entries can carry private notes to preserve moderation context, and org blocks now show who applied the block and when it expires. The blocked-user search flow has also been separated from the 'block a new user' action to reduce accidental blocks.

Review who applied an org block and when it expires, to audit moderation actions across a large organization.

> 📍 In Organization settings, go to Moderation › Blocked users, then filter by block reason (e.g., `Harassment`) and review the applied-by and expiry details for each entry.

— Gives exact UI path and filter values, no API surface.

Better tools for managing blocked users…

WAS THIS USEFUL? 👍 👎

JFrog

Feedback ⓘ 1 RELEASE · 2026-08-21 BLOG ↗

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

> JFrog Artifactory added On-Behalf-Of (OBO) OAuth 2.0 token exchange support for Amazon Bedrock AgentCore Gateway, letting AI agent requests reach Artifactory authenticated as the originating end user instead of a shared service account.

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

01 On-Behalf-Of identity propagation for AgentCore Gateway NEW

86

Artifactory Access now supports On-Behalf-Of (OBO) token exchange (RFC 8693) with Amazon Bedrock AgentCore Gateway, so requests from AI agents calling Artifactory's MCP server as a registered gateway target are exchanged for short-lived, per-user JFrog tokens rather than a shared credential — preserving audit fidelity, least privilege, and containment. Configuration is done via `onBehalfOfTokenExchangeConfig` on the gateway target, and the new `listingMode` field (requires a recent AWS CLI) enables dynamic tool listing so tool discovery also runs under the user's token. Identity mapping relies on JFrog Access OIDC integration: JFrog trusts the external IdP issuer (`$IDP_ISSUER`) and maps verified access-token claims (e.g. `preferred_username` , `team`) to JFrog users, groups, or project roles via the JFrog OIDC integration API, requiring the token's `aud` claim to equal the AgentCore Gateway URL (`$GATEWAY_URL`) and `iss` to match `$IDP_ISSUER` . This capability requires JFrog EnterpriseX or higher. The post also compares this approach against header whitelisting (forwarding `X-Authenticated-User` `-style` headers) and three-legged OAuth (authorization-code grant with a per-user token vault in AgentCore Identity), noting OBO is the only option delivering verifiable per-user identity at Artifactory without a second login.

.
.
.

— Names exact config keys, claims and API but no runnable command shown.

Propagating User Identity From AI Agent…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ 1 RELEASE · 2026-08-21 NOTES ↗

Infisical is the open-source platform for secrets, certificates, and privileged access management.

> Infisical v0.162.23 expands the certificate manager with PKCS#12 keystore import and raises the PKI post-sync command limit to support longer automation scripts.

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

01 **PKCS#12 keystore import in certificate manager** NEW

48

The certificate manager now supports importing PKCS#12 keystores, adding a new certificate import format alongside existing options.

— Names the format but not import steps or UI path v0.162.23

02 **PKI post-sync command limit raised to 8192 characters**

IMPROVED

45

The character limit for PKI post-sync commands was raised to 8192 characters, enabling longer post-sync scripts to be configured.

— Gives exact new limit but no config path or before value v0.162.23

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Bunkerity BunkerWeb

Feedback ⓘ 1 RELEASE · 2026-08-21

NOTES ↗

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

> BunkerWeb 1.6.14 centralizes mTLS CA/CRL distribution via the Scheduler, adds TOTP-based 2FA and pinned server-side sessions, hardens default headers and Permissions-Policy, and extends Redis-backed metrics alongside per-service CrowdSec controls and a UI certificate validator.

└──▷ WHAT SHIPPED · 11 FEATURES most completely described first ? what's the number?

01 Redis-backed metrics sync and blocked-request retention NEW

90

Adds `METRICS_REDIS_TTL` (default `2592000` seconds) to control expiry of Redis metrics keys, with active keys refreshed each sync; the metrics plugin now syncs blocked-request data to Redis for a centralized view across instances, buffering new reports per-worker and syncing once Redis memory frees to avoid loss, with atomic list and facet updates. `METRICS_MAX_BLOCKED_REQUESTS_REDIS` now honors `k` / `m` suffixes when bounding report reads. Without Redis, the report cap is the buffer size multiplied by NGINX workers per instance and is cleared on restart unless Redis persistence is enabled; long-term retention requires forwarding logs to syslog or a SIEM.

Set Redis metrics TTL to 30 days so inactive blocked-request data ages out automatically while active data stays live.

```
$ METRICS_REDIS_TTL=2592000
```

— Names TTL default, config keys and OOM buffering behavior v1.6.14

02 TOTP-based two-factor authentication for admin accounts NEW

90

Adds TOTP-based 2FA with secrets stored encrypted in the database using keys kept in `/var/lib/bunkerweb/.totp_encryption_keys.json`. The `TOTP_ENCRYPTION_KEYS` environment variable (space-separated keys or a JSON map, each key exactly 43 characters) supplies or rotates keys without relying on auto-generation, with old keys kept alongside new ones during rotation so existing enrollments stay valid. Mounting a volume on `/data` in the `bunkerweb-ui` container persists TOTP encryption keys, Biscuit keys, and the Flask secret across recreation; at startup the UI resolves keys from the file, then `TOTP_ENCRYPTION_KEYS`, then

generates a random set (invalidating all 2FA enrollments) if neither is usable, and an invalid key value (e.g. fewer than 43 characters) is discarded and replaced by a random one, leaving 2FA unusable until cleared manually.

Supply your own TOTP encryption key so 2FA enrollments survive container recreations even without a persistent volume.

```yaml
# TOTP_ENCRYPTION_KEYS must be exactly 43 characters (e.g. output of passlib generate_secret())
TOTP_ENCRYPTION_KEYS: "aB3dEfGhIjKlMnOpQrStUvWxYz0123456789AbCdE"
```

Rotate TOTP encryption keys while keeping existing 2FA enrollments valid by supplying both the old and new key in a JSON map.

```yaml
TOTP_ENCRYPTION_KEYS: '{"new": "aB3dEfGhIjKlMnOpQrStUvWxYz0123456789AbCdE", "old": "zY9xWvUtSrQpOnMlKjIhGfEdCbA9876543210ZyXw"}'
```

Persist TOTP encryption keys and Biscuit keys across container recreations by mounting a volume on `/data` in the UI service.

```yaml
services:
  bw-ui:
    image: bunkerity/bunkerweb-ui:1.6.14
    volumes:
      - bw-ui-data:/data   # persists Flask secret, TOTP encryption keys, and Biscuit keys

volumes:
  bw-ui-data:
```

— Names env var, exact key length, file path and rotation flow `product docs`

---

03 **mTLS CA bundle and CRL distribution via Scheduler** NEW

85

Adds `MTLS_CA_CERTIFICATE_DATA` and `MTLS_CA_CERTIFICATE_PRIORITY` to supply the trusted client CA bundle inline as base64 or PEM (e.g. via the Web UI) instead of mounting files, and `MTLS_CRL_DATA` / `MTLS_CRL_PRIORITY` to supply a PEM-encoded certificate revocation list the same way, enabling revoked-certificate rejection. The Scheduler validates, caches under `/var/cache/bunkerweb`, and distributes both to every BunkerWeb instance automatically on each job run with no per-instance volume

mounts; if the cache volume is not persistent, mTLS stays disabled after a scheduler restart until the first job run redistributes the bundle.

Enable mTLS by pointing the Scheduler at a CA bundle file and a CRL — no per-instance mounting needed.

```
$ MTLS_CA_CERTIFICATE_PRIORITY=file
MTLS_CA_CERTIFICATE_DATA=
MTLS_CRL_PRIORITY=file
MTLS_CRL_DATA=
```

Supply the CA bundle and CRL inline (e.g. in a containerised deployment where file mounts are inconvenient).

```
$ MTLS_CA_CERTIFICATE_PRIORITY=data
MTLS_CA_CERTIFICATE_DATA=<base64-or-PEM-bundle>
MTLS_CRL_PRIORITY=data
MTLS_CRL_DATA=<base64-or-PEM-crl>
```

Supply a client CA bundle inline so the scheduler distributes it to all instances without mounting files into each container.

```yaml
MTLS_CA_CERTIFICATE_DATA: "<base64-or-PEM-encoded-CA-bundle>"
MTLS_CRL_DATA: "<base64-or-PEM-encoded-CRL>"
```

— Names exact config keys, priority sources and cache path product docs

04 **Inbound X-SSL-* header stripping** BREAKING 75

Incoming `X-SSL-*` request headers (e.g. `X-SSL-Client-Verify: SUCCESS`) are now stripped on every site under HTTP/1.1, HTTP/2 and HTTP/3 regardless of whether mTLS is enabled, so only values derived from the verified TLS handshake reach the application. Deployments behind an upstream proxy that sets these headers must capture the value via a `server-http` Nginx `set` directive and re-inject it through `REVERSE_PROXY_HEADERS`, since `proxy_set_header` alone runs too late.

— Explains mechanism and gives exact migration workaround product docs

05 **Per-service CrowdSec settings and offline mode** NEW 75

Every `CROWDSEC_*` setting is now `multisite`, letting individual services point at different Local API or AppSec endpoints; setting `CROWDSEC_API` to an empty string on

a service retains AppSec inspection without the decision lookup. Adds `DISABLE_ONLINE_API` to the `all-in-one` image (matching the official CrowdSec image name) to run CrowdSec without registering to the Central API.

Run the all-in-one container in an air-gapped environment where CrowdSec must not call home to the Central API.

```
$ docker run -e DISABLE_ONLINE_API=true bunkerity/bunkerweb-all-in-one:1.6.14
```

— Names settings, empty-string trick, and a runnable command v1.6.14

06 **Permissions-Policy blocks Client Hints and Chrome AI APIs**

70

BREAKING

The default `PERMISSIONS_POLICY` now blocks all `ch-*` Client Hints directives, breaking services relying on `Accept-CH` or `Sec-CH-Prefers-Color-Scheme` until the policy is adjusted to re-allow specific hints (e.g. ch-dpr=(self)); it also now denies Chrome's built-in AI APIs by default.

Re-allow device-pixel-ratio Client Hints after the new default Permissions-Policy blocks all ch-* directives, so responsive-image pipelines keep working.

yaml

```
PERMISSIONS_POLICY: "accelerometer=(), ..., ch-dpr=(self)"
```

Re-allow a specific client hint after the new default blocks all ch-* directives, so adaptive-image services continue to receive DPR information.

```
$ # In your BunkerWeb service environment or UI, append to Permissions-Policy:
ch-dpr=(self)
```

— Names the header and directives with a concrete fix v1.6.14

07 **Server-side sessions pinned to IP and User-Agent**

BREAKING

65

Sessions are pinned to client IP and User-Agent with a default idle lifetime of 12 hours and a hard cap of 7 days; cookies can be made permanent without a separate long-lived token, still bound by those limits, and are revoked on logout, password change, or `Wipe other sessions`. Existing long-lived `remember me` tokens are now rejected and deleted on the next request, logging out users who relied on year-long persistence;

`SESSION_LIFETIME_HOURS` and `SESSION_ABSOLUTE_HOURS` can be extended to replicate long-lived sessions.

— Names lifetime defaults and the two override variables v1.6.14

THINNER COVERAGE BELOW

08 **Certificate validation button in UI** NEW 50

Adds a Validate certificate button on the custom certificate settings page that checks a pasted or uploaded key pair for mismatches, encrypted keys, expiry, and missing server name coverage before the service is saved.

— Names UI element and checks but no config surface v1.6.14

09 **KEEP_CONFIG_ON_RESTART moved out of web UI** BREAKING 50

`KEEP_CONFIG_ON_RESTART` is no longer configurable via the web UI or database; it must be set in the environment or in `/etc/bunkerweb/variables.env` on Linux, since the entrypoint reads it before the database is reachable.

— Names exact config file path and reason for the move v1.6.14

10 **Biscuit token single-block enforcement** BREAKING 30

The API now rejects Biscuit tokens carrying more than a single block; clients that appended extra blocks to their tokens will have requests refused.

— States the change but gives no remediation path v1.6.14

11 **QUERY HTTP method allowed by default** IMPROVED 30

`ALLOWED_METHODS` and bundled service templates now allow the `QUERY` HTTP method by default.

— One-line default change with no further mechanism given v1.6.14

└──▷ **ALSO FROM THESE RELEASES**

Supply a mTLS client CA bundle and CRL entirely inline to the Scheduler so no volume mounts are needed on BunkerWeb instances.

```
$ MTLS_CA_CERTIFICATE_DATA=<base64-or-pem-bundle>
MTLS_CRL_DATA=<base64-or-pem-crl>
```

└──▷ **BREAKING ON UPGRADE**

- ! The default Permissions-Policy now blocks all `ch-*` (Client Hints) directives; services relying on `Accept-CH` or `Sec-CH-Prefers-Color-Scheme` will lose Client Hints support until the policy is adjusted to re-allow the required `ch-*` features.
- ! Incoming `X-SSL-*` headers are stripped before reaching the application on every site; setups behind an upstream proxy that sets these headers must capture the value in a `server-http` Nginx `set` directive (e.g. `set $trusted_ssl_verify $http_x_ssl_client_verify`) and re-inject it via `REVERSE_PROXY_HEADERS` — using `proxy_set_header` alone is insufficient because the header is already empty by the time it is evaluated.
- ! When mTLS is enabled and the scheduler cache at `/var/cache/bunkerweb` is not persistent, mTLS remains disabled after a scheduler restart until the first job run completes and redistributes the CA bundle; a persistent cache volume is required for strict enforcement across restarts.
- ! Recreating the `bunkerweb-ui` container without a persistent volume on `/data` wipes TOTP encryption keys; every user loses their 2FA enrollment and must re-enroll.
- ! Setting `TOTP_ENCRYPTION_KEYS` to an invalid value (e.g. fewer than 43 characters, such as the placeholder `changeme`) causes the key to be discarded and replaced by a random one, leaving 2FA unusable until the variable is cleared manually — unlike leaving the variable unset, this also prevents the admin enrollment from being reset automatically.
- ! `KEEP_CONFIG_ON_RESTART` is no longer listed as a configurable setting in the web UI or database — it must be set in the environment or in `/etc/bunkerweb/variables.env` on Linux, as it is read by the entrypoint before the database is reachable.
- ! Existing `remember me` long-lived tokens are rejected and deleted on the next request; users relying on year-long browser persistence will be logged out and must re-authenticate. Extend `SESSION_LIFETIME_HOURS` and `SESSION_ABSOLUTE_HOURS` to replicate long-lived sessions.
- ! Biscuit tokens carrying more than a single block are now rejected by the API; any client that appended extra blocks to its tokens will have its requests refused.

WAS THIS USEFUL? 👍 👎

DataDome

Feedback

ⓘ

1 RELEASE · 2026-08-21

BLOG ↗

DataDome is a bot management and API protection platform that defends against automated attacks, web scraping, and credential abuse.

> DataDome's Envoy module gained GraphQL support in version 2.8.0, with configurable settings to enable and extract GraphQL operation details.

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

01 **GraphQL support in Envoy module** NEW 86

Envoy module 2.8.0 adds GraphQL request handling via the `ENABLE_GRAPHQL_SUPPORT` setting (set to `true` to enable) and the `GRAPHQL_ENDPOINT` setting (default: `/graphql`), which specifies one or more comma-separated endpoints from which the GraphQL operation name and type are extracted.

— Names exact config keys, default value, and mechanism. Envoy module 2.8.0 (2026-08-21)

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

NETRESEC

Feedback ⓘ 1 RELEASE · 2026-08-21 BLOG ↗

NETRESEC provides network security tools and services for detecting, analyzing, and investigating cyber threats and malware.

> NETRESEC's FlowCarp gained a new detection model for the CNCMachineRMS RAT's C2 protocol, backed by a free public demo endpoint and detailed IOC/evasion documentation for defenders.

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

01 **CNCMachineRMS C2 detection in FlowCarp** NEW 85

Adds a statistical detection model to FlowCarp for the CNCMachineRMS binary C2 protocol, identifying sessions on TCP port 443 that mimic HTTPS but are actually unencrypted, without needing TLS decryption. A free public demo server at `https://demo.flowcarp.com` accepts raw PCAP uploads via `curl --data-binary @suspicious.pcap https://demo.flowcarp.com` and returns a JSON alert if CNCMachineRMS traffic is present.

Quickly triage a suspicious PCAP from an incident or sandbox run to check for CNCMachineRMS C2 traffic without standing up any local tooling.

```
$ curl --data-binary @suspicious.pcap https://demo.flowcarp.com
```

Test a suspicious PCAP for CNCMachineRMS C2 traffic using the free FlowCarp demo server — useful for quick triage without deploying local tooling.

```
$ curl --data-binary @suspicious.pcap https://demo.flowcarp.com
```

— Runnable curl command against a named public endpoint.

CNCMachineRMS C2 Protocol (2026-08-21)

02 **CNCMachineRMS IOCs and DoH evasion documented** NEW 70

Publishes network IOCs for CNCMachineRMS: four confirmed C2 server IPs on port 443 (`89.110.110.119:443`, `85.158.110.78:443`, `89.124.79.98:443`, `144.124.242.171:443`), two C2 domains (`notepadreleased[.]com`, `triotmelon[.]com`), and four file hashes for known samples. Also documents the malware's DNS-over-HTTPS evasion technique, routing lookups to `dns.google`,

`cloudflare-dns.com`, and `dns.quad9.net` over TLS before the port-443 C2 connection, so local DNS logs won't capture the domain lookup and defenders should instead alert on DoH activity preceding port 443 C2 traffic.

.

— Concrete IOCs given but no direct tool action beyond ingestion.

CNCMachineRMS C2 Protocol (2026-08-21)

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Alibaba Cloud OpenSandbox

Feedback ⓘ

1 RELEASE · 2026-08-21 NOTES ↗

Secure, Fast, and Extensible Sandbox runtime for AI agents.

> OpenSandbox's egress sidecar now auto-allows its own OTLP metrics export destination, preventing the deny-all policy from blocking observability traffic.

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

01 Auto-allow rule for sidecar OTLP metrics egress NEW

85

The egress sidecar reads `OTEL_EXPORTER_OTLP_METRICS_ENDPOINT` (preferred) or `OTEL_EXPORTER_OTLP_ENDPOINT` at startup and automatically injects the destination into the deny-all policy — domains via the DNS proxy, IPs into the static nft allow sets. This always-allow rule survives policy updates and reloads, though `deny.always` still takes precedence over it.

Point the egress sidecar at your OTLP collector so metric export is never blocked by the deny-all policy, without writing a manual allow rule.

```
$ OTEL_EXPORTER_OTLP_METRICS_ENDPOINT=https://otel-
collector.internal:4318/v1/metrics
```

— Names exact env vars, mechanism, and precedence rule. docker/egress/v1.1.7

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

kata-containers

Feedback ⓘ 1 RELEASE · 2026-08-21 NOTES ↗

Kata Containers is an open source project and community working to build a standard implementation of lightweight Virtual Machines (VMs) that feel and perform like containers, but provide the workload isolation and security advantages of VMs. https://katacontainers.io/

> Kata Containers 4.1.0 adds new hypervisor options (OpenVMM and Cloud Hypervisor VM templates), NVSwitch GPU passthrough, TDX guest compatibility improvements in dragonball, and a broad set of kata-deploy configuration, security and release-packaging hardening.

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

01 **kata-deploy configuration, security and Helm improvements**

IMPROVED 80

kata-deploy gains numerous configuration and hardening additions: `shared_fs=none` to skip shipping virtiofsd when it is not needed, EROFS backing mode configuration support, Helm chart values overlay support, `podLabels` support on per-node Jobs, TEE key advertisement rendering from the Helm chart, job-mode dispatcher confinement to trusted nodes, isolation of guest debug within runtime classes, removal of `nsenter` with tightened host privileges, binding of nydus-snapshotter to the node's actual CRI unit, and Helm-driven rollout when mounted ConfigMaps change.

— Names every config key and flag but no usage walkthrough 4.1.0

02 **Release packaging and QEMU build hardening** IMPROVED

65

The static release tarball is now split by runtime into per-runtime tarballs, a `qemu-no-shared-fs` build variant is added, confidential-guest QEMU builds are trimmed, and QEMU features and devices are disabled by default to reduce the attack surface.

— Names build variant and packaging changes with rationale 4.1.0

03 **TDX guest support improvements in dragonball** IMPROVED

60

dragonball improves TDX guest compatibility by adding the `VIRTIO_F_ACCESS_PLATFORM` flag for TDX virtio devices and a hypercall handler (`VcpuExit::Hypercall`) for TDX.

— Names both internal mechanisms added for TDX 4.1.0

THINNER COVERAGE BELOW

04 **OpenVMM and Cloud Hypervisor VM template support in runtime-rs**

NEW 45

runtime-rs adds support for the OpenVMM hypervisor and for Cloud Hypervisor VM templates, expanding the hypervisor and VM provisioning options available to Kata sandboxes.

– Names both hypervisors, no configuration detail given 4.1.0

05 **NVSwitch device passthrough support** NEW 45

Adds `nvswitch-class` device passthrough support in the runtime, enabling NVSwitch-based GPU topologies to be exposed to guests.

– Names the device class but not setup steps 4.1.0

06 **listxattr added to dragonball seccomp allowlist** IMPROVED 45

Adds `listxattr` to the dragonball seccomp allowlist in runtime-rs, enabling workloads that call this syscall to run without being blocked.

– Names exact syscall added, no broader scope given 4.1.0

07 **L3 forwarding network model in runtime-rs** NEW 40

Adds an L3 forwarding network model for runtime-rs, providing an additional network topology option for Kata sandboxes.

– States capability but not how to select it 4.1.0

08 **New guest extension support in runtime-rs** NEW 40

runtime-rs adds a generic devkit debug guest extension and composable VM extensions for IBM SEL, extending the guest extension framework.

– Names two extensions but not their usage 4.1.0

09 **Local CUDA repository support for NVIDIA GPU extensions** NEW 35

Adds local CUDA repository support for the NVIDIA GPU (`nvgpu`) extensions.

– Bare mention of feature, no setup detail 4.1.0

10 **PCI utilities extracted into pcilibs-rs crate** IMPROVED 30

Extracts PCI utilities into a standalone `pcilibs-rs` crate.

– Internal refactor named but no user-facing impact described 4.1.0

11 **Shared block storage source mapping in runtime-rs** IMPROVED

30

runtime-rs now shares block storage source mapping to reduce duplication.

— States benefit but no mechanism or scope 4.1.0

12 **Synchronous sandbox cgroup placement in runtime-rs** IMPROVED

30

The runtime now places itself in the sandbox cgroup synchronously in runtime-rs, improving resource accounting.

— States benefit but lacks technical detail 4.1.0

13 **virtiofsd bumped to 1.14.0** IMPROVED

25

Bumps the bundled virtiofsd dependency to version 1.14.0.

— Version bump only, no changelog of what changed 4.1.0

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

Tenzir

Feedback ⓘ 1 RELEASE · 2026-08-22 BLOG ↗

Tenzir is a data pipeline platform that ingests, transforms, and analyzes security telemetry and logs for threat detection.

> Tenzir's v6.13.0 release adds native OpenTelemetry ingestion and node-wide pipeline parallelism, alongside parallelized file-reading, a reworked `parallel` block, and a flat-naming convention for built-in operators and functions.

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

01 **Native OTLP ingestion via `accept_otlp`** NEW 93

Adds the `accept_otlp` operator to receive OpenTelemetry logs, metrics, and traces over OTLP/HTTP or OTLP/gRPC, supporting JSON and binary Protobuf, gzip compression, TLS and mTLS, signal filtering, and bounded request sizes; set `schema="record"` for direct attribute-key access.

Ingest OpenTelemetry signals from an OTEL Collector or SDK over OTLP/HTTP on port 4318

```
$ accept_otlp "0.0.0.0:4318", transport="http"
```

— Full transport, format, and security options named plus runnable command

Tenzir Node v6.13.0: Parallel pipelines…

02 **Node-wide pipeline parallelism controls** NEW 93

Adds the `tenzir.parallelism` configuration option (also settable via the `TENZIR_PARALLELISM` environment variable) to configure parallelism node-wide for all pipelines that don't specify their own degree. A `// parallelism:` directive on a pipeline takes precedence over the `--parallelism` flag, which takes precedence over `tenzir.parallelism`; use `// parallelism: disabled` to opt a single pipeline out.

Set node-wide parallelism so every pipeline without an explicit directive runs on up to 8 cores

```yaml
tenzir:
  parallelism: max,limit_partitions=8
```

— Config key, env var, precedence order and opt-out flag all named

Tenzir Node v6.13.0: Parallel pipelines…

03 **Parallel `summarize` with periodic emission** IMPROVED 73

The `summarize` operator can now run on multiple cores even when `emit: <duration>` is set, enabling parallel periodic aggregation; use `// parallelism: 1` or the `window` operator when aligned emission intervals are required. It also delivers faster `summarize` performance on interleaved group keys.

— Names exact flags and fallback operator, mechanism partly explained

Tenzir Node v6.13.0: Parallel pipelines…

04 **`decimal_format="float"` option for `read_parquet`** NEW 72

Adds a `decimal_format="float"` option to the `read_parquet` operator to read Arrow `Decimal128` values as potentially lossy real values; default behavior reads them as strings.

— Exact flag and default behavior named, easy to apply product docs

05 **Parallel file-reading across storage sources** IMPROVED 71

`from_file`, `from_s3`, `from_google_cloud_storage`, and `from_azure_blob_storage` now run as multiple instances when the pipeline is parallelized, splitting discovered files across instances so large directories are read on multiple cores simultaneously.

— Names all four operators but no config example given

Tenzir Node v6.13.0: Parallel pipelines…

06 **Reworked `parallel` block semantics** IMPROVED 67

The `parallel` block is now semantically transparent: operators inside observe the same events as they would without it, `route_by` is automatically inferred and deprecated as an explicit option, and wrapping a source in `parallel N { … }` no longer produces N copies of every event.

— Explains before/after behavior but no runnable example

Tenzir Node v6.13.0: Parallel pipelines…

07 **`fused=none` parallelism channel semantics** IMPROVED 65

Setting `fused=none` in a parallelism directive (e.g. `// parallelism: 6,fused=none`) now consistently means plain buffered channels throughout rather than restricting to a small buffer, and sequential pipelines now run without fused channels.

— Exact directive syntax given but behavior change is niche

Tenzir Node v6.13.0: Parallel pipelines…

08 Flat names for built-in operators and functions DEPRECATED

65

Renames all built-in module-qualified operators and functions to flat names, e.g. `context::enrich` becomes `context_enrich`, `ocsf::cast` becomes `ocsf_cast`, and `ai::prompt` becomes `ai_prompt`; old spellings emit a deprecation warning and continue working until the next major release.

— Concrete rename examples but no migration command shown

Tenzir Node v6.13.0: Parallel pipelines…

THINNER COVERAGE BELOW

09 Multiple pipeline arguments in tree-sitter-tql NEW

50

tree-sitter-tql v1.2.0 adds support for multiple pipeline arguments separated by commas in operators like `fork_merge`, letting them accept more than one `{ … }` pipeline block without a syntax error.

— Names the operator and version but no usage example product docs

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ 1 RELEASE · 2026-08-21 NOTES ↗

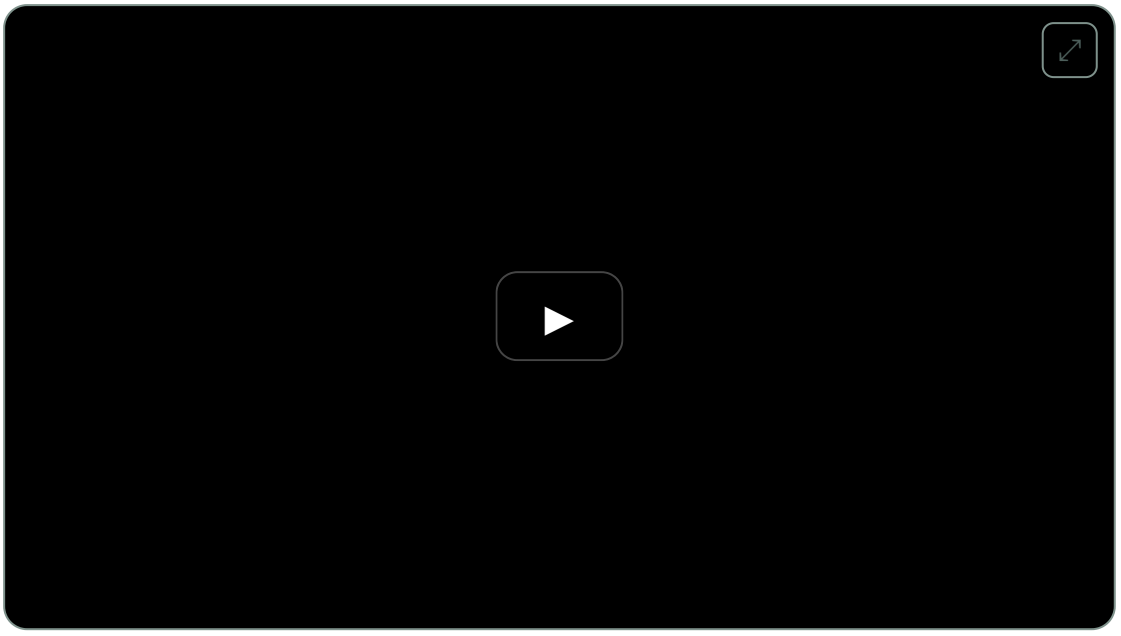
> Iru shipped AI-guided update guidance in the Windows Update Library, a new System Tray app for employee update visibility, and new Auto Apps for Windows and Mac.

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

01 AI integration in Windows Update Library Item NEW

40

Adds native Iru AI integration to the Windows Update Library Item to provide a guided update experience for users.



– Names the surface but no mechanism detail on how AI guides users.

What's New in Iru | August 21st

02 System Tray app for update visibility NEW

35

A new Windows System Tray app gives employees visibility into pending updates along with their deadlines.

– Describes purpose but not how deadlines are set or displayed.

What's New in Iru | August 21st

03 New Auto Apps for Windows and Mac NEW

20

New Auto Apps were added for both Windows and Mac platforms.

– No detail on what the Auto Apps actually do. What's New in Iru | August 21st

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta

Feedback ⓘ 1 RELEASE · SEEN 2026-08-22 NOTES ↗

Immuta is a data governance and access control platform that enables secure, policy-driven data sharing while maintaining compliance and privacy.

> Immuta adds summary metrics to the Access Requests page for governors and refines Databricks Unity Catalog ingestion to skip inaccessible catalogs upfront.

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

01 **Unity Catalog ingestion skips inaccessible catalogs** IMPROVED 67

Databricks Unity Catalog crawl logic now filters out, at ingestion time, catalogs where the Immuta workspace has read-only access and isolated catalogs where the Immuta service principal lacks `MANAGE` permission, rather than failing later during policy enforcement.

— Names exact permission and mechanism change but no command to run snapshot-20260822

THINNER COVERAGE BELOW

02 **Access request queue summary tiles** NEW 58

Three summary tiles added to the Access Requests page show pending request count, average age of pending requests, and stewards with at least one pending request assigned. Hovering over a tile reveals a breakdown of top individuals or groups by request volume.

— Describes exact tiles and hover behavior but no config surface snapshot-20260822

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

ConductorOne

Feedback ⓘ 1 RELEASE · 2026-08-21 BLOG ↗

ConductorOne is an identity governance platform that automates access management, provisioning, and compliance workflows across enterprise applications.

> ConductorOne's biggest release this window is a native MCP gateway integration for Google Gemini Enterprise with per-user OAuth attribution, alongside a new Copilot Studio MCP client, smart revocation that traces access through nested groups and role hierarchies, a major c1i CLI expansion (with a breaking URL-flag change), reusable/local policies, and a round of public API and review-workflow improvements.

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

01 **Gemini Enterprise MCP gateway integration** NEW 100

A new ConductorOne MCP gateway integration for Google Gemini Enterprise routes tool calls through `https://<your-tenant>-mcp.conductor.one/v1`, attributing each call to the individual user's identity rather than a shared service account. It publishes a pre-registered OAuth client metadata document at `https://<your-tenant>.conductor.one/auth/v1/client-metadata/gemini-enterprise` (the document URL is the `client_id`, using PKCE with no client secret), and exposes `authorization_endpoint` and `token_endpoint` discoverable via `https://<tenant>-mcp.conductor.one/.well-known/oauth-authorization-server`. Setup requires overriding the Google Cloud org policy `constraints/discoveryengine.managed.disableCustomMcpServerConnector` to Off, enabling six GCP APIs (`discoveryengine.googleapis.com`, `connectors.googleapis.com`, `integrations.googleapis.com`, `secretmanager.googleapis.com`, `iam.googleapis.com`, `orgpolicy.googleapis.com`), optionally allowlisting egress via `constraints/discoveryengine.allowedEgressFqdns`, setting the connector's Authorization URL Parameters to `&resource=https://<your-tenant>-mcp.conductor.one/v1`, and using space-separated scopes `openid profile email offline_access`. Connections are visible under `AI access management > AI clients` (admin) or profile menu `AI & API > AI connections` (user), and calls are gated by C1 toolsets and access profiles.

Check and disable the org policy that blocks custom MCP data connectors before creating the Gemini Enterprise data store.

```
$ # Check current effective policy
gcloud org-policies describe \
  discoveryengine.managed.disableCustomMcpServerConnector \
  --project=YOUR_PROJECT_ID --effective

# Save override to policy.yaml, then apply
cat > policy.yaml <<'EOF'
name:
projects/YOUR_PROJECT_ID/policies/discoveryengine.managed.disable
CustomMcpServerConnector
spec:
  rules:
    - enforce: false
EOF
gcloud org-policies set-policy policy.yaml
```

Enable all six required Google Cloud APIs for the Gemini Enterprise MCP integration in one command.

```
$ gcloud services enable \
  discoveryengine.googleapis.com \
  connectors.googleapis.com \
  integrations.googleapis.com \
  secretmanager.googleapis.com \
  iam.googleapis.com \
  orgpolicy.googleapis.com \
  --project=YOUR_PROJECT_ID
```

Confirm the C1 OAuth client metadata and discover authorization/token endpoints for your tenant before configuring the Gemini Enterprise data store.

```
$ # Confirm the client metadata document resolves
curl -s https://YOUR_TENANT.conductor.one/auth/v1/client-
metadata/gemini-enterprise

# Discover authorization_endpoint and token_endpoint
curl -s https://YOUR_TENANT-mcp.conductor.one/.well-known/oauth-
authorization-server
```

— Full endpoint, config-key and runnable gcloud/curl detail provided.

August 21, 2026 (2026-08-21)

02 **c1i CLI expanded with app, auth and policy commands** NEW

90

c1i CLI v0.5.0 adds `apps create`, `apps delete`, `apps set-owners`, `auth token`, `mcp gateway list-tools`, `mcp gateway call`, and a `policies` command family, rounding out API coverage for agent-focused automation. List commands now emit properly-typed Booleans, counts, and nullable fields instead of stringified or inconsistent JSON, and all responses map onto a single predictable exit-code scheme.

List available MCP gateway tools via the c1i CLI to verify which governed tools are callable through the gateway.

```
$ c1i mcp gateway list-tools --server-url
https://<tenant>.conductor.one
```

Call a governed MCP tool through the c1i CLI to verify the gateway is reachable and the tool executes under the caller's identity.

```
$ c1i mcp gateway call --server-url
https://<tenant>.conductor.one/mcp <tool-name> '{"arg": "value"}'
```

— Exact new commands with runnable examples. August 21, 2026 (2026-08-21)

---

03 **Smart revocation via Revoke inherited access** NEW

75

A new Revoke inherited access toggle on an app's details page traces access upstream through direct grants, nested groups, role hierarchies, and delegated provisioning and removes it, instead of stopping at the first result. The toggle is off by default and shows a downstream-effects banner before an approver commits.

— Exact toggle name and default state, UI path given. August 21, 2026 (2026-08-21)

---

04 **c1i tenant URL and mcp gateway flag change** BREAKING

70

`c1i` now requires a full `https://` tenant URL — bare short-name shortcuts are rejected, and a non-HTTPS URL is rejected instead of silently upgraded. The `mcp gateway` command's `--url` flag has been renamed to `--server-url`, breaking any script or alias that still passes `--url` to `mcp gateway`.

— Exact flag rename and requirement given as migration note. August 21, 2026 (2026-08-21)

---

05 **Grant source scoping on automation revoke steps** NEW

70

A new `Grant source` setting (`Direct` or `All`) on automation revoke steps lets automations scope revocations to only direct grants, avoiding collateral removal of group-inherited access.

Scope an automation's revoke step to direct grants only, so group-inherited access is left intact during automated cleanup.

> 📍 In the automation editor, open the revoke step, find the `Grant source` setting, and select `Direct` .

Scope an automation revoke step to direct grants only, so group-inherited access is not collaterally removed during cleanup.

> 📍 In the automation editor, open the revoke step, then set Grant source to `Direct` instead of `All` .

— Named setting with exact steps to configure it. August 21, 2026 (2026-08-21)

---

06 **Public API managed application support** NEW 65

The public API can now list connector-discovered applications, inspect their managed state, and promote an unmanaged application to managed — returning HTTP `409` on a repeat promotion and exposing the managed application ID in the binding response.

— Concrete status codes and response fields, no exact path given.

August 21, 2026 (2026-08-21)

---

07 **Push rule execution history** NEW 65

Push-rule execution history is now available at `Directory sources > Executions` (all connectors) or via `Show execution history` in a push rule's row menu — each entry records who was pushed, when, which attributes changed, and the rejection reason on failure.

— Exact UI navigation and logged fields named. August 21, 2026 (2026-08-21)

---

08 **Local and reusable policies** NEW 60

Policies can now be created local to a single app or entitlement via inline creation alongside the global option on each configuration slot. A policy rule or baseline can also reference another policy of the same type, letting a shared flow (e.g. a manager-then-owner approval chain) be defined once and reused — updates to the shared policy propagate automatically to every policy that references it.

— Describes mechanism and propagation but no UI path given. August 21, 2026 (2026-08-21)

THINNER COVERAGE BELOW

09 **Custom attribute columns in access reviews** NEW 55

Access review table columns can now include any reviewer-visible app user attribute (job title, department, cost center) — up to 10 per campaign — so reviewers no longer need to hover each row.

— Concrete limit and example attributes named. August 21, 2026 (2026-08-21)

10 **Task page detail drawer and request filtering** IMPROVED 55

Task reviewers can open a read-only details drawer for the entitlement or resource under review directly from the task page, with no special role required. An app's Requests tab can now be filtered to `Request` or `Revoke` tasks, and the `Open revocations` stat card on the app's Overview tab links directly to that filtered view.

— Two related task-page UI additions named with exact filter values.

August 21, 2026 (2026-08-21)

11 **Copilot Studio as MCP gateway client** NEW 45

Copilot Studio agents can now call IT-approved tools through the C1 MCP gateway, with every call authorized and logged under the calling user's own identity, and can be published to Microsoft 365 Copilot and Teams.

— Named integration described only in prose, no config surface given.

August 21, 2026 (2026-08-21)

12 **MCP servers in global search** NEW 45

MCP servers are now searchable from global search (`Cmd+K`), with recently opened servers surfaced under `Your recent searches` .

— Exact UI entry point named, but a small feature. August 21, 2026 (2026-08-21)

13 **Grant history records get unique IDs** IMPROVED 40

Each record returned by the public API's past-grants search now includes a unique `id` field, making individual grant history entries directly addressable.

— Single named field, no usage example given. August 21, 2026 (2026-08-21)

14 **OCSF time field in circuit-breaker logs** IMPROVED 40

Automation circuit-breaker system-log events now include the standard OCSF `time` field, enabling direct SIEM ingestion without a custom parsing workaround.

— Single named field change, thin description. August 21, 2026 (2026-08-21)

15 **Permission-hidden info toast** IMPROVED 30

Pages that hide information due to role restrictions now display a toast — 'Some information on this page is hidden based on your permissions' — instead of showing blank fields with no indication.

— Cosmetic change, exact toast text given but no config surface.

August 21, 2026 (2026-08-21)

16 **Task and review table column sizing** IMPROVED 30

Task lists and access review tables now size columns to fit content and scroll horizontally instead of compressing all columns to the same width; long entitlement names truncate with an ellipsis and reveal the full name on hover.

— UI polish described without any named surface. August 21, 2026 (2026-08-21)

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! `c1i` v0.5.0 now requires a full `https://` tenant URL — bare short-name shortcuts are rejected, and a non-HTTPS URL is rejected instead of silently upgraded.

! The `mcp gateway` command's `--url` flag has been renamed to `--server-url`; any scripts or aliases using `--url` with `mcp gateway` will break.

! `c1i` now requires a full `https://` tenant URL — bare short-name shortcuts are rejected, and a non-HTTPS URL is rejected instead of silently upgraded.

! The `mcp gateway` command's `--url` flag is renamed to `--server-url`; any script passing `--url` to `mcp gateway` must be updated.

WAS THIS USEFUL? 👍 👎

SPIFFE spire

Feedback ⓘ 1 RELEASE · 2026-08-21 NOTES ↗

The SPIFFE Runtime Environment

> SPIRE v1.15.3 adds a Slurm workload attestor and an Azure Blob Storage bundle publisher, alongside a batch of new agent and attestor configuration options and a debug CLI subcommand.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 Disable Workload API and SDS API on agent endpoint NEW

75

New `disable_workload_api` and `disable_sds_api` agent options disable the Workload API and SDS APIs on the public endpoint, useful for reducing attack surface where those APIs aren't needed.

Disable the Workload API and SDS API on a public-facing agent endpoint to reduce attack surface in environments where those APIs are not needed.

```hcl
disable_workload_api = true
disable_sds_api = true
```

— Two flags with exact config file snippet shown v1.15.3

02 debug getinfo CLI subcommand NEW

70

New `debug getinfo` subcommand in the `spire-server` and `spire-agent` CLIs provides access to the Debug APIs.

Inspect live agent state via the Debug API without needing to query the gRPC endpoint directly.

```
$ spire-agent debug getinfo
```

— Exact runnable command shown for both CLIs v1.15.3

03 Kubernetes workload attestor options: kubelet client and namespace labels

NEW

65

The `k8s` workload attestor gains `disable_kubelet_client` to turn off use of the kubelet client, and an opt-in `enable_namespace_labels` option that produces `ns-label` selectors from namespace labels.

Enable namespace label selectors in the Kubernetes workload attestor to write SPIFFE policies based on namespace-level labels.

```hcl
enable_namespace_labels = true
```

— Both flags named with example config for one v1.15.3

04 **Initial trust bundle from Workload API endpoint** NEW 60

New `trust_bundle_spiffe_workload_api` agent configuration option lets an agent fetch its initial trust bundle from a Workload API endpoint, simplifying nested agent deployments.

— Config key and purpose named, no example of exact usage v1.15.3

05 **Pod-UID-based agent IDs in k8s_psat attestor** NEW 60

New `use_pod_uid_for_agent_id` option in the `k8s_psat` node attestor derives agent IDs from pod UIDs instead of node UIDs.

— Named option with before/after behaviour, no runnable example v1.15.3

06 **File-based account list for aws_iid attestor** NEW 60

New `account_list_file` option in the `aws_iid` node attestor lets `verify_organization` 's account list be sourced from a file instead of the AWS Organizations API.

— Named option and mechanism, no config example given v1.15.3

THINNER COVERAGE BELOW

07 **Azure Blob Storage bundle publisher** NEW 50

New `azure_blob` BundlePublisher plugin publishes the trust bundle to Azure Blob Storage.

— Plugin named, no config or usage detail given v1.15.3

08 **WIT-SVID support in svid.v1 API client** NEW 50

Adds WIT-SVID building support in the `svid.v1` API client and marshalling support for WIT-SVID keys.

— Names API and SVID type, no usage example v1.15.3

09 **Broker API in agent/server health checks** IMPROVED 35

The Broker API endpoint is now included in health checks when enabled.

— Brief mention with no configuration or scope detail v1.15.3

10 **Slurm workload attestor** NEW 30

New Slurm workload attestor plugin.

— Only the plugin name is given, no mechanism v1.15.3

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-21 BLOG ↗

Clerk is an authentication and user management platform that handles sign-up, login, and identity verification for applications.

> Clerk shipped a new CLI for zero-config authentication setup across a dozen frameworks, a shadcn/ui theme package, and custom OAuth scopes for MCP client access control.

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

01 Clerk CLI for scaffolding and managing auth setup NEW

95

Adds `npx clerk@latest init`, which auto-detects the framework and package manager, installs the SDK, and wires up the provider, middleware, auth routes, and env file for Next.js, React, React Router, Nuxt, TanStack Start, Astro, Vue, JavaScript/Vite, Expo, Express, Fastify, iOS, and Android; it supports `--framework <framework>` and `--pm <package-manager>` to override auto-detection in empty directories, and `--app <application_id>` to link an existing Clerk application. When the framework supports it, `init` writes temporary development keys to the project env file so no Clerk account is required upfront, and the app is claimed later via `npx clerk@latest auth login`. Also adds `npx clerk@latest doctor` to verify the auth setup and report configuration issues, `npx clerk@latest apps list --json` to list available Clerk applications with names and IDs, and `npx clerk@latest enable orgs` to turn on Organizations (multi-tenancy, team invitations, roles and permissions, enterprise SSO) on an existing project.

Scaffold a new project in an empty directory with a specific framework and package manager, then verify the setup is correct.

```
$ npx clerk@latest init --framework react --pm pnpm && npx clerk@latest doctor
```

— Names every subcommand and flag with a runnable example. product docs

02 shadcn/ui theme via @clerk/ui package NEW

75

Adds the `@clerk/ui` package for shadcn/ui theme integration: apply it via `appearance={{ theme: shadcn }}` on `ClerkProvider` and `@import '@clerk/ui/themes/shadcn.css'` in global CSS.

Apply the shadcn/ui theme to Clerk components when the project already uses shadcn/ui.

```javascript
import { shadcn } from '@clerk/ui/themes'
// in your provider:
<ClerkProvider appearance={{ theme: shadcn }}>{children}
</ClerkProvider>
```

— Exact prop and import path with code example given. product docs

THINNER COVERAGE BELOW

03 **Custom OAuth scopes for MCP clients** NEW 25

Supports custom OAuth scopes, enabling precise and discoverable access control for MCP clients connecting to your API.

— No mechanism, config keys, or example given. Custom OAuth scopes (2026-08-21)

WAS THIS USEFUL?  

Okta Identity Engine

Feedback

1 RELEASE · 2026-08-19

BLOG ↗

Okta Identity Engine is a modern identity platform providing authentication, authorization, and user management for applications.

> Okta Identity Engine made unified AI agent identity generally available in Production and added Cross App Access (XAA) support letting AI agents and apps connect without runtime user consent, alongside new admin controls for agent identity conflicts and a breaking change to the Delegation workflow.

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

01 **Unified identity for AI agents (GA)** NEW — 65

Okta now resolves AI agent identities across every registration path into a single authoritative record per org, with each agent carrying an `externalId` and `platform` value for uniqueness. This capability is now generally available (GA) in Production, enabling identity lifecycle management for non-human AI agent principals.

— Explains mechanism and fields but no API call shown

Weekly release 2026.08.2 (2026-08-19)

02 **Cross App Access (XAA) for AI agents** NEW — 65

Cross App Access (XAA) enables connections between custom SSO-agentic requesting apps and SSO resource apps across OIDC and SAML SSO protocols, removing the need for user consent at runtime. An AI agent can be bound with an OIDC SSO app so they share the same credentials, with admin-controlled visibility over which actions the agent can take on behalf of a user.

— Clear mechanism and protocol scope, no exact config shown product docs

03 **User access tab replaces Delegation tab** BREAKING — 65

A new User access tab configures direct user authentication for AI agents, replacing the Delegation tab workflow. Okta for AI Agent orgs that previously used the Delegation tab to configure AI agent access through delegation links must reconfigure them using the User access tab.

— Exact UI tabs named, clear migration path product docs

THINNER COVERAGE BELOW

04 **AI Agent Providers admin endpoints** NEW — 50

New endpoints support listing, reviewing, and resolving staged identity conflicts, and managing provider priority for AI agents, documented under AI Agent Providers.

— Endpoints named only generically, no paths given product docs

└──▷ **BREAKING ON UPGRADE**

! Okta for AI Agent orgs that previously used the Delegation tab to configure AI agent access through delegation links must reconfigure them using the User access tab.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Pomerium

Feedback ⓘ 3 RELEASES · 2026-08-21 NOTES ↗

Pomerium is an identity-aware proxy that provides secure, zero-trust access to internal applications and services.

> Pomerium's biggest work this window is MCP OAuth server-side hardening — CIMD-based client identification, a per-cluster Dynamic Client Registration toggle, prefixed access/refresh tokens, RFC 8252-compliant loopback matching, and relaxed client_id validation for ChatGPT — alongside new Enterprise Console startup flags, distroless container images, and stable GitHub node_id keying for directory users.

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

01 **Per-cluster toggle for inbound Dynamic Client Registration**

NEW 70

Adds a runtime toggle in Pomerium Zero cluster settings to enable inbound Dynamic Client Registration (DCR) per cluster; the feature ships off by default rather than relying on a single global default. Enable it in the Zero dashboard by going to Cluster Settings for the target cluster and toggling Inbound Dynamic Client Registration on.

Enable inbound Dynamic Client Registration for a specific cluster so downstream MCP clients can self-register, without enabling it globally.

📍 In the Zero dashboard, go to Cluster Settings for the target cluster and toggle Inbound Dynamic Client Registration on.

— Exact UI navigation to enable the toggle is given MCP OAuth Compatibility & Hardening

02 **OAuth Client ID Metadata Document support for MCP** NEW

60

Pomerium adds OAuth Client ID Metadata Document (CIMD) support for inbound MCP routes, letting downstream clients such as Claude Code, ChatGPT, and MCP Inspector identify themselves via the lighter-weight CIMD registration path instead of full Dynamic Client Registration.

— Names mechanism and clients but gives no config step

MCP OAuth Compatibility & Hardening

03 **Console startup flags for directory sync and route setup** NEW

60

The Enterprise Console adds `-directory-provider` and related startup flags so directory sync can be fully configured at launch without using the web UI, plus a `-url` startup flag that lets the console automatically create its own Pomerium route on first boot, eliminating a manual setup step for new installs.

— Names both flags but omits exact syntax or values Enterprise Console Updates

04 **Distroless base image for container hardening** IMPROVED 60

Pomerium container images now use the `base-nossl-debian12` distroless variant instead of `base-debian12`, structurally removing `libssl3` from every deployment.

— Names exact image variants and removed library Platform & Image Hardening

THINNER COVERAGE BELOW

05 **RFC 8252-compliant loopback redirect matching** IMPROVED 55

Loopback redirect URI matching is now RFC 8252 §7.3-compliant: port differences on localhost/loopback URIs are ignored, allowing CLI clients like Claude Code to use a different loopback port than the one originally registered.

— Explains exact behavior change, automatic with no configuration

MCP OAuth Compatibility & Hardening

06 **Stable GitHub node_id keying for directory users** IMPROVED 55

The GitHub directory provider can now use GitHub's stable `node_id` as the primary key for directory users instead of the GitHub login (username), so policies survive account renames or username reuse.

— Explains mechanism and benefit but no config key given Platform & Image Hardening

07 **Prefixed MCP access and refresh tokens** NEW 50

MCP access and refresh tokens are now issued with prefixes `pom_mat_…` and `pom_mrt_…`, making Pomerium-issued credentials identifiable in logs and catchable by secret-scanning tools.

— Exact token prefixes named; change is automatic, no action needed

MCP OAuth Compatibility & Hardening

08 **Relaxed client_id metadata URL validation for ChatGPT**

IMPROVED 50

Relaxes `client_id` metadata URL validation to accept URLs containing a query string, enabling compatibility with ChatGPT's MCP connector, which was previously rejected with a 400.

— Names exact field and prior error code, no user action needed

MCP OAuth Compatibility & Hardening

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NVIDIA OpenShell

Feedback ⓘ 1 RELEASE · 2026-08-21 NOTES ↗

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

> OpenShell v0.0.111 focuses on sandbox networking controls and observability, adding transparent TCP egress with DNS-based policy enforcement, OTLP tracing for the Podman driver, and new standalone compute drivers.

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

01 **Network egress and DNS policy controls for sandboxes** NEW

46

Adds transparent TCP egress and policy DNS enforcement for Docker and Podman runtimes, backed by a new policy DNS correlation store that ties network policy decisions to DNS lookups in Docker and Podman sandboxes, and establishes a direct TCP egress foundation for policy-governed outbound connections.

– Describes mechanism and scope but no config keys or flags given v0.0.111

02 **OTLP traces for Podman driver** NEW

38

Exports Podman driver traces over OTLP, giving observability into sandbox compute activity.

– Names OTLP export but no endpoint or config detail v0.0.111

03 **Standalone first-party compute drivers** NEW

17

Adds standalone first-party compute drivers for sandbox execution.

– Bare description, no mechanism or scope given v0.0.111

04 **Canonical main process for sandbox runtime** IMPROVED

16

Adds a canonical main process to the sandbox runtime.

– Bare mention with no further explanation v0.0.111

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue