THE CYBER TOOLCHAIN

NO. 024

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 24, 2026 · EVERY WEEKDAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 25 tools

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Promptfoo ⓘ snapshot-20260724 NOTES ↗

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

└──▷ TRY IT

View promptfoo log output directly from the CLI for debugging rate limiting or provider issues.

```
$ promptfoo logs
```

- Adds xAI Voice provider overrides: apiBaseUrl, websocketUrl, and function call support.
- Adds test-level prompts filter, providers filter, and per-test structured output configuration at the individual test case level.
- Adds `--share` and `--no-share` flags to Model Audit for controlling cloud sharing.
- Adds multi-input variable support to red team scans for testing systems with complex input structures.
- Adds early scan stop when plugins fail to generate test cases.

+1 more

WAS THIS USEFUL? 👍 👎

Portkey ⓘ snapshot-20260724 NOTES ↗

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

Portkey adds vault-backed Secret References, weekly rate limits, cross-cloud GCS log delivery, and four new guardrail integrations.

└──▷ USE IT

Apply Zscaler Detection Policies to all LLM traffic through Portkey so your existing Zscaler policy class governs both inputs and outputs.

json

```
Add Zscaler AI Guard as a guardrail with your policyId:

{
  "guardrails": {
    "beforeRequestHook": {
      "provider": "zscaler",
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    },
    "afterRequestHook": {
      "provider": "zscaler",
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    }
  }
}
```

- Adds Secret References so virtual keys and integrations can pull credentials at runtime from AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault via `secret_mappings`, keeping secrets out of Portkey config.
- Supports weekly budget and usage windows (`rpw`) for rate-limit policies so spend caps align with weekly planning cycles.
- Adds endpoint-scoped rate limits so chat completions, embeddings, and admin routes can each carry independent limits instead of a single global cap.
- Enables GCS log storage from AWS-hosted gateways via GCP Workload Identity Federation (`GCP_WIF_AUDIENCE` , `GCP_WIF_SERVICE_ACCOUNT_EMAIL`), eliminating long-lived cross-cloud credentials.
- Lets organization admins include archived workspaces in analytics graphs and filters, with workspace slugs now accepted alongside IDs in saved filters.

+11 more

WAS THIS USEFUL? 👍 👎

Veracode ⓘ SCA updates (2026-07-24) BLOG ↗

Veracode is a cloud-based application security platform that helps developers identify and remediate vulnerabilities in their code through static, dynamic, and software composition analysis.

Veracode SCA gains malicious package detection, Poetry support, SBOM generation, new agent flags, and Unified Policy for workspaces.

└──▷ **TRY IT**

Capture structured JSON scan output for automation while keeping a readable summary visible in the terminal during a CI run.

```
$ srcclr scan . --json=results.json --show-cli
```

Assign a Unified Policy to an SCA workspace so builds break based on policy criteria — navigate to the workspace policy settings in the Veracode Platform.

> 📍 In the Veracode Platform, go to SCA › Workspaces › <your workspace> › Policies, then select Assign Policy and choose or set a default Unified Policy.

- Adds `--appname` flag to the SCA agent scan command (Commercial region) to link scan results directly to an application profile, requiring OAuth or HMAC authentication.
- Adds `--show-cli` flag to the SCA agent scan command: used with `--json`, it saves structured JSON output to a file while also displaying a human-readable summary in the console.
- Enables Unified Policy assignment to SCA agent-based scan workspaces in the US Federal region, allowing policies to create issues and break builds based on defined criteria.
- Enables Unified Policy assignment to SCA agent-based scan workspaces in the European region.
- Adds ability to configure the Vulnerability Severity policy rule (Commercial region) to break builds only when an upgrade path to fix a vulnerable component is available.

+14 more

└──▷ **BREAKING ON UPGRADE**

! The default SCA agent download URL changes from `https://download.sourceclear.com` to `https://sca-downloads.veracode.com`; the legacy URL will be retired after 12 months.

! The default backend API URL for Commercial region customers changes from `https://api.sourceclear.io` to `https://sca-api.veracode.com`; the legacy URL will be retired after 12 months.

! The default backend API URLs for European region customers change from `https://platform-backend.analysiscenter.veracode.eu` to `https://sca-api.veracode.eu`; legacy URLs will be retired after 12 months.

! The default backend API URLs for US Federal region customers change from `https://platform-backend.analysiscenter.veracode.us` to `https://sca-api.veracode.us`; legacy URLs will be retired after 12 months.

WAS THIS USEFUL? 👍 👎

Cloudsmith ⓘ

`CLI v1.19.0: Authenticate automatically with OIDC and Docker credential discovery (2026-07-23)`

BLOG ↗

Cloudsmith is a cloud-native, universal package management platform that enables teams to securely host, manage, and distribute software artifacts across multiple formats.

Cloudsmith CLI v1.19.0 adds automatic OIDC credential discovery for CI/CD and a Docker credential helper for registry auth.

› Automatically discovers OpenID Connect (OIDC) credentials for all major CI/CD platforms, eliminating the need for static API keys in pipelines.

› New Docker credential helper automatically authenticates to Cloudsmith registries without manual `docker login` steps.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Self-Hosted ⓘ 2026.7 - Required (2026-07-24)

BLOG ↗

GitGuardian Self-Hosted is an on-premise security tool that detects and prevents secrets leakage in Git repositories.

GitGuardian Self-Hosted 2026.7 adds Public Monitoring, native Gerrit scanning, Smart Notifiers, Gateway API routing, and 9 new secret detectors.

└──▷ USE IT

Enable the Public Monitoring worker in your Helm values to start processing public perimeter secret-leak incidents.

```yaml
celeryWorkers:
  public-monitoring:
    enabled: true
```

Deploy the GitGuardian MCP server in your cluster (beta) by enabling it in your Helm values.

```yaml
mcpServer:
  enabled: true
```

- Adds Public Monitoring on self-hosted instances (license-gated): monitor your public perimeter for secret leaks and route public incidents through custom webhooks, Slack, Microsoft Teams, ServiceNow ticketing, and severity-filtered email alerts via the new `celeryWorkers.public-monitoring` worker.
- Adds native Gerrit VCS source with full historical scanning, real-time detection via webhook plugin, and read-replica support for cloning.
- Adds Smart Notifiers: define per-channel rules to filter alerts, tickets, and emails by severity, validity, secret type, or GitGuardian tag across internal and public monitoring.
- GitHub check runs now scan oversized pull requests by cloning the repository branch instead of calling the GitHub API commit-by-commit, and can skip merge commits to protect shared API rate limits; thresholds and behavior are configurable from the admin area GitHub preferences.
- Adds Gateway API routing to direct external traffic through the Kubernetes Gateway API instead of the legacy Ingress resource, plus support for HPA or KEDA autoscaling for webapps.

+5 more

└──▷ BREAKING ON UPGRADE

- ! Helm installations: `logCollector.collectMigrationLogs` is now enabled by default and injects a native sidecar requiring Kubernetes 1.29 or later; on older clusters you must set `logCollector.collectMigrationLogs: false`.
- ! Endpoint Protection (Developer Endpoint Protection) is not functional on self-hosted instances in this release — do not upgrade if you rely on Endpoints without contacting GitGuardian support first.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Prowler ⓘ 5.36.0 NOTES ↗ CODE ↗

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

Prowler 5.36.0 adds SageMaker secrets detection, grouped Jira dispatch, cleaner Attack Paths UI, and regionless OCI onboarding.

└──▷ GET THIS VERSION

```
$ git clone --branch 5.36.0 https://github.com/prowler-cloud/prowler.git
# already have the repo? check out this version:
$ git checkout 5.36.0
```

- New `sagemaker_notebook_instance_no_secrets` check scans SageMaker notebook instance `OnCreate` and `OnStart` lifecycle scripts for hardcoded API keys, passwords, tokens, and connection strings.
- Grouped Jira workflows (Prowler Cloud/Private Cloud): send selected findings or Finding Groups to Jira as one grouped issue or separate issues, with deep links and filter context preserved.
- Attack Paths query selector now hides confirmed-empty queries for the selected scan, surfacing only paths that returned data; full catalog published on Prowler Hub.
- OCI provider setup no longer requires a region — SDK handles region discovery automatically, simplifying onboarding.
- New AI agents tutorials and overview banner guide users through connecting AI agents to Prowler Cloud via the API to query and act on security posture programmatically.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Docker Sandboxes ⓘ v0.37.0 NOTES ↗

Docker Sandboxes isolates AI model execution environments to safely test and run untrusted code with controlled resource access.

Docker Sandboxes v0.37.0 adds SSH access to sandboxes, shared agent skills, port publishing, and richer audit records.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.37.0 https://github.com/docker/sbx-releases.git
# already have the repo? check out this version:
$ git checkout v0.37.0
```

└──▷ TRY IT

Set up SSH access once so every sandbox becomes reachable by name from any OpenSSH-compatible client or remote-dev tool.

```
$ sbx setup ssh
```

Connect to a running sandbox for an interactive shell or one-shot command without leaving your terminal.

```
$ ssh my-sandbox.sbx
```

Import host agent skills into the shared store (preview first, then apply) so new sandboxes automatically inherit them.

```
$ sbx skills import --dry-run && sbx skills import --force
```

Start a sandbox with a published port and without shared skills, e.g. for an isolated web-service test.

```
$ sbx run -p 8080:8080 --no-share-skills my-image
```

- Adds SSH access to sandboxes (experimental): run `sbx setup ssh` once to register a managed `*.sbx` SSH config entry, then connect to any sandbox with `ssh <name>.sbx` from OpenSSH clients or remote-dev tools like VS Code and Cursor.
- Adds `sbx skills import` to discover and import host agent skills into a persistent store shared across sandboxes, with `--dry-run` to preview and `--force` to overwrite existing skills.
- New `--no-share-skills` flag on `sbx run` / `sbx create` lets individual sandboxes opt out of the shared skills store.
- Adds `-p` / `--publish` flag to `sbx create` and `sbx run` for publishing sandbox ports at creation time.
- New `DOCKER_SANDBOXES_PROXY=system` env var routes sandbox egress through the host OS proxy (macOS/Windows), including PAC auto-config URLs.

`+1 more`

WAS THIS USEFUL? 👍 👎

Tencent Cloud CubeSandbox ⓘ v0.6.0 NOTES ↗

Instant, Concurrent, Secure & Lightweight Sandbox for AI Agents.

CubeSandbox v0.6.0 adds Kubernetes deployment, E2B-compatible volumes, template aliases, configurable Host headers, and node isolation.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.6.0
https://github.com/TencentCloud/CubeSandbox.git
# already have the repo? check out this version:
$ git checkout v0.6.0
```

└──▷ TRY IT

Isolate a compute node before maintenance so no new sandboxes are scheduled to it while existing ones keep running.

```
$ cubemastercli node isolate <node-id>
```

Lift isolation on a node after maintenance is complete to allow new sandbox scheduling to resume.

```
$ cubemastercli node unisolate <node-id>
```

- Deploy CubeSandbox control-plane and compute nodes on TKE, standard Kubernetes, or k3s via Helm Chart, with support for rolling upgrades and canary releases.
- New E2B-compatible Volume framework lets operators plug in custom backend storage via Binary or RPC plugins, with Create/Destroy/Attach/Detach lifecycle hooks and sandbox–Volume binding APIs.
- Template aliases can be set at creation time, allowing sandboxes to be launched by alias name instead of template ID.
- New `network.maskRequestHost` parameter lets operators configure the `Host` header forwarded into a sandbox (e.g. `localhost:${PORT}`), improving compatibility with Host-based routing inside sandboxes.
- New compute node isolation commands (`cubemastercli node isolate` / `unisolate`) let operators drain a node for maintenance without disturbing running sandboxes.

+9 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Corellium ⓘ Release 7.12.0 (2026-07-23) BLOG ↗

Corellium is a virtualization platform enabling security researchers and developers to run iOS, Android, and Linux systems virtually for testing and analysis.

Corellium 7.12 adds Metal graphics support, KASAN/HWASAN kernel debugging, PongoOS pre-boot, MCP Server integration, and folder downloads.

└──▷ HOW TO FIND IT

Download an entire app data directory from a virtual device at once for forensic analysis or artifact collection, instead of extracting files one by one.

> 📍 In the Corellium console, open the Files Browser for your virtual device, navigate to the target folder (e.g., `/var` /mobile/Containers/Data/Application/<app-uuid>), select the folder, and click Download to retrieve the full directory tree as a compressed archive.

- Adds initial Apple Metal graphics support via a software-based implementation targeting Flutter applications, enabling apps that previously failed to launch or render in virtualized iOS environments.
- Adds Android KASAN/HWASAN kernel debugging support, giving security researchers visibility into memory corruption, use-after-free, and buffer overflow issues at the kernel level.
- Enables pre-boot of supported Apple devices directly into PongoOS before iOS loads, streamlining jailbreak research and low-level boot customization workflows.
- Introduces an integrated Model Context Protocol (MCP) Server as the first phase of next-generation API/SDK modernization, allowing AI tools and autonomous agents to interact directly with Corellium environments.
- The Files Browser now supports downloading entire folders and nested sub-folder structures from virtual devices in a single operation.

+2 more

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Elastic Detection Rules ⓘ dev-v2.0.0 NOTES ↗

Detection Rules are automated security patterns that identify suspicious activities and threats by analyzing event data in real-time.

Elastic Detection Rules dev-v2.0.0 adds MITRE ATT&CK v19 support and 80+ new rules across AWS, Azure, Linux, GCP, and network protocols.

└──▷ GET THIS VERSION

```
$ git clone --branch dev-v2.0.0 https://github.com/elastic/detection-rules.git
# already have the repo? check out this version:
$ git checkout dev-v2.0.0
```

- Adds MITRE ATT&CK v19 threat mapping support, automatically applied on Elastic Stack >= 9.5 with dynamic tag assignment.
- New AWS rules: CloudTrail management event disabling, GuardDuty suppression detection, IAM console login without MFA, IAM user login from multiple geolocations, account closure, org-leave attempts, Cognito unauthenticated identity pool credential issuance, Bedrock AgentCore execution role misuse, Bedrock prompt/completion credential detection, ECR public access grants, SageMaker suspicious lifecycle config, Lambda cross-account invocation, Lambda high-frequency invocation, Lambda event source mapping creation, Lambda layer external sharing, Lambda function URL public access, Lambda deletion, Lambda policy cross-account update, Backup vault deletion/lock removal, Backup recovery point deletion, KMS imported key material deletion, IAM inline policy added to group, IAM permissions boundary modification, IAM login profile creation/modification, and IAM account password policy deletion.
- New Azure/Entra ID rules: AKS API server proxying requests to Kubelet, ROPC authentication with unknown client ID, VM extension installation and suspicious child process via CustomScript, VM configuration modification, VM Managed Run Command with unusual principal, VM Extension CRUD from unusual ASN, VM Boot Diagnostics retrieval, Serial Console connect with unusual user/ASN, Azure AD Graph access with unusual user/ASN, unusual client/user, suspicious user-agent, and 4xx error surge, Azure AD Graph potential enumeration via ROADrecon, Entra ID AiTM phishing-kit chain, device registration with phishing-kit default OS build, conditional access MFA bypass via Microsoft Graph, multiple device registrations by single user, impossible travel sign-in, ROPC auth/legacy client detection, OAuth device code sign-in to Azure AD Graph enumeration, OAuth device code phishing, TAP creation, guest-to-member promotion, and OAuth redirect URI abuse.

› New GCP/GKE rules: GKE Kubernetes rule conversions (parts 1 and 2) adding coverage for GKE API and workload activity.

› New Linux rules: kernel core_pattern container escape, systemd service override config file creation, Redis CONFIG SET SSH authorized key injection, potential Redis Lua activity, AMSI bypass via RPC NdrClientCall hook (PowerShell 4104), Linux LPE rule migrations from Endpoint Rules to Detection Rules (parts 1 and 2).

+6 more

└──▷ **BREAKING ON UPGRADE**

! Emitted related integrations now use `>=` instead of `^` for version constraints, which is expected to not function on Elastic stacks older than 9.5 — ensure your stack schema map and/or packages.yml match your stack version before upgrading.

! MITRE ATT&CK v19 mappings are automatically applied on stack >= 9.5; to use v19 on stacks <= 9.4 you must explicitly set `DR_THREAT_MAPPING_FRAMEWORK="MITRE ATT&CK"` and `DR_THREAT_MAPPING_VERSION=1`, or add `threat_mapping_framework: "MITRE ATT&CK"` and `threat_mapping_version: "19"` to your _config.yaml.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero ⓘ runZero 5.0.260723.0 (2026-07-23) BLOG ↗

runZero is a continuous asset intelligence platform that discovers, inventories, and monitors all connected devices and software across networks in real-time.

runZero 5.0.260723.0 adds Purdue level visibility for Dragos-enriched OT/ICS assets.

› Displays Purdue Model levels for assets sourced or enriched from Dragos integration, enabling OT/ICS network segmentation visibility.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter ⓘ changelog · Jul 24, 2026 CHANGELOG ↗

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

DNSFilter redesigns the Query Log with multi-select filters, cross-org switching, and URL-persistent views, plus live local domain sync in the macOS client.

└──▷ HOW TO FIND IT

Investigate DNS activity across multiple organizations simultaneously without losing your filter context — select multiple Sites and Roaming Clients in the filter bar, then switch orgs to compare.

> 📍 In the console, go to Query Log › top-level filter bar, select multiple Sites/Roaming Clients/Relays/Users, then use the org toggle to switch organizations while filters remain active.

Share a specific DNS investigation view with a teammate by bookmarking the URL after applying filters — active filters persist in the URL automatically.

> 📍 In the console, go to Query Log, apply the desired filters and timeline via the custom timeline picker, then copy and share the resulting URL.

Set the Query Log as your default landing page so it opens immediately on login.

> 📍 In the console, navigate to the Query Log page, then use the 'Set as Default Landing Page' option to bookmark it as your homepage.

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

+3 more

└──▷ BREAKING ON UPGRADE

! Starting with macOS Roaming Client v2.4.4, macOS 13 Ventura or later is required; macOS 12 Monterey is no longer supported.

WAS THIS USEFUL? 👍 👎

Zscaler Private Access ⓘ changelog · Jul 24, 2026 BLOG ↗

Zscaler Private Access provides zero-trust network access to internal applications without exposing them to the public internet.

ZPA gains Autonomous App Shield and Autonomous User-to-App Segmentation to continuously protect and isolate private apps.

- New Autonomous App Shield continuously assesses private applications via App Connectors, reasons about per-app risk using global threat intelligence, and applies targeted protections automatically — adapting when developers ship new releases without manual re-tuning.
- New Autonomous User-to-App Segmentation narrows user and third-party access to exactly the applications required, reducing lateral movement opportunity when a foothold is gained.
- New Zscaler Lateral Threat Bundle packages application hiding (ZPA), autonomous segmentation, and App Shield into a unified offer for layered private-app defense.

WAS THIS USEFUL? 👍 👎

Slack Nebula ⓘ v1.11.0 NOTES ↗ CODE ↗

A scalable overlay networking tool with a focus on performance, simplicity and security

Nebula v1.11.0 adds macOS network-change rebinding, IPv6 rejects, stdin/stdout cert I/O, and six breaking config changes.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.11.0 https://github.com/slackhq/nebula.git
# already have the repo? check out this version:
$ git checkout v1.11.0
```

└──▷ USE IT

Keep macOS laptops connected when they roam between wifi and wired — disable the default rebind behaviour only if your environment requires it.

```yaml
# In config.yml — omit or set to true to enable (the default)
listen:
  rebind_on_network_change: true
```

- Adds automatic listener rebind and lighthouse re-query on macOS when the underlay network changes (wifi↔wired), controlled by `listen.rebind_on_network_change` (default `true`).
- Generates IPv6 reject packets for blocked traffic, matching existing IPv4 reject behavior.
- Accepts `-` as a path in `nebula-cert` subcommands to read from stdin or write to stdout.
- Searches for both `config.yml` and `config.yaml` automatically in service and CLI modes.
- Adds version labels to Docker/OCI images.

+8 more

└──▷ BREAKING ON UPGRADE

! Logging has switched from logrus to Go's structured `slog`: log levels are now upper case (`level=INFO`), trace prints as `level=DEBUG-4`, timestamps are always RFC3339Nano, `logging.timestamp_format` is ignored, and some

messages were reworded. Embedders must update constructors to accept `*slog.Logger`.

- ! `firewall.inbound_action` and `firewall.outbound_action` were each applied to the opposite direction; this is now corrected. If you set either, you are currently getting the behavior of the other and likely need to swap them before upgrading.
- ! On Windows, Nebula now installs WFP PERMIT filters for the nebula adapter and listener port by default, bypassing Windows Defender Firewall (WDF) inbound rules for those. Set `tun.windows_bypass_wdf` and `listen.windows_bypass_wdf` to `false` to leave WDF in charge.
- ! On Windows, the Nebula device network category is now set to `private` instead of whatever Windows chose (usually `Public`). Set `tun.network_category` to `unset` to keep the old behavior.
- ! Reject packets for non-TCP traffic now use ICMP code 13 (communication administratively prohibited) instead of code 3 (port unreachable). Anything keying off the old code needs updating.
- ! The SSH debug server's profiling commands (`start-cpu-profile`, `save-heap-profile`, `save-mutex-profile`) are now confined to `sshd.sandbox_dir` (default `$TMP/nebula-debug`). Absolute paths outside that directory are rejected; the directory is not created automatically. Scripts using paths elsewhere must set `sshd.sandbox_dir` accordingly.

WAS THIS USEFUL? 👍 👎

◆ **Identity & Access Management**

iDenfy ⓘ changelog · Jul 24, 2026 CHANGELOG ↗

iDenfy is an identity verification platform that uses AI-powered document scanning and biometric facial recognition to verify user identities and prevent fraud.

- › Supports shared identity verification across companies in KYB (Know Your Business) workflows.

WAS THIS USEFUL? 👍 👎

Clerk ⓘ changelog · Jul 24, 2026 CHANGELOG ↗

Clerk is a user authentication and identity management platform that provides developers with pre-built UI components, APIs, and SDKs for sign-up, sign-in, and user management.

Clerk CLI gains a `clerk mcp` command group for one-command MCP server registration across 10 AI clients, plus configurable default OAuth scopes.

└──▷ TRY IT

Register the Clerk MCP server across all detected AI clients in one shot — useful for onboarding a new machine or CI environment.

```
$ clerk mcp install --all
```

Target only Claude Code and Cursor for MCP registration when you don't want to touch other clients on the machine.

```
$ clerk mcp install --client claude --client cursor
```

Verify your Clerk MCP server is reachable (not just installed) after registration, to catch config or network issues early.

```
$ clerk doctor
```

- Adds `clerk mcp install` command that detects supported AI clients (Claude Code, Cursor, GitHub Copilot, Windsurf, Gemini CLI, Codex, opencode, OpenClaw, Warp, Hermes Agent) and registers the Clerk MCP server in each, user-globally, in a single interactive step.
- Adds `--client` and `--all` flags to `clerk mcp install` for non-interactive, scripted registration targeting specific or all supported clients.
- Adds `clerk mcp run`, a built-in stdio↔Streamable-HTTP bridge that replaces the `npx mcp-remote` dependency; updating the CLI updates the bridge everywhere with no re-install.
- Adds `clerk mcp list` to display every Clerk MCP entry registered across all clients, and `clerk mcp uninstall` to remove them selectively.
- Extends `clerk doctor` to probe the configured Clerk MCP server with a real `initialize` handshake, distinguishing 'not installed', 'installed but unreachable', and 'working' states.

+1 more

└──▷ **BREAKING ON UPGRADE**

! Beginning January 18, 2027, CBC-mode SSL/TLS cipher suites (including `TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA`, `TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA`, `TLS_RSA_WITH_AES_256_CBC_SHA`, and others) will be removed from Clerk-managed subdomains (`clerk.example.com`, `accounts.example.com`); clients that can only negotiate a CBC cipher will fail the TLS handshake after their subdomain's certificate rotates.

WAS THIS USEFUL? 👍 👎

Ping Identity ⓘ

changelog · Jul 24, 2026 CHANGELOG ↗

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

Ping Identity AIC adds email scripting, secrets in nodes, Rich Authorization Requests, SAML metadata import options, and a new validate-secrets endpoint.

- Adds `emailService` binding for next-generation scripts, enabling email delivery directly from journey scripts.
- New realm-level OAuth2 Provider setting enables Rich Authorization Requests (RFC 9396) processing when an RCS is configured, and controls advertisement of `authorization_details_types_supported` on the well-known endpoint.
- Individual OAuth 2.0 clients can now be configured with `Authorization Detail Type(s)` to restrict which authorization detail types each client may request.
- Dynamic client registration requests now support `authorization_details_types` to declare the authorization detail types a client supports at registration time.
- SAML metadata import now supports skipping signature verification.

+12 more

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Immuta ⓘ snapshot-20260724 NOTES ↗

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

└──▷ HOW TO FIND IT

When onboarding a data product with sensitive columns, configure distinct request forms so access requestors and masking-exception requestors each see only relevant questions and agreements.

> 📍 In the Immuta UI, navigate to your data product › Settings › Request Forms, then configure a separate form under Data Access Requests and another under Masking Exception Requests, adding tailored questions and data use agreements to each.

Write a subscription policy that restricts AI-restricted data to human users only, excluding all agent accounts system-wide.

> 📍 In the Immuta UI, create a subscription policy targeting your data source, set the condition to allow only members of the built-in group 'All Users' (excluding 'All Agents'), and save.

Enable OAuth 2.0 for the Alation integration to avoid recurring token expiry issues.

> 📍 In the Immuta UI, navigate to Settings › Integrations › Alation, select OAuth 2.0 as the authentication method, and provide the required OAuth credentials.

- Adds separate request forms for data products: one for data access requests and one for masking exception requests, each with independent questions and data use agreements.
- Expands Snowflake manual policy resync to cover both data policies and subscription policies simultaneously.
- Adds automated sensitive data discovery (regex, dictionary, and column-name identifiers) for the Amazon Redshift viewless integration.
- Adds OAuth 2.0 authentication support to the Alation integration, eliminating repeated token refresh cycles required by API key authentication.
- Embeds Immuta documentation and an AI-powered chatbot directly in the Immuta UI via a new Documentation icon in the top menu.

+1 more

└──▷ BREAKING ON UPGRADE

- ! Snowflake integration architecture update (rolling out July 23, 2026): Immuta policy role names used to grant users access to Snowflake data sources will be changed and recreated; any automation or tooling that references specific policy role names by their current values will need to be updated.
- ! Immuta will no longer display integration error banners when integration validation fails for Snowflake; errors will only be visible through the connections experience.
- ! When a data source is disabled in Immuta, Immuta now revokes data owners' Immuta-managed grants in the data platform (behavior change releasing 8/1 by default, 9/1 for all).

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant ⓘ powerbi-v1.0.1 NOTES ↗

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 150+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

CISO Assistant gains a native Power BI connector for direct BI integration.

└──▷ GET THIS VERSION

```
$ git clone --branch powerbi-v1.0.1 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout powerbi-v1.0.1
```

› New native Power BI connector enables direct integration between CISO Assistant and Power BI for security reporting and dashboards.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro ⓘ V2 to V3 Agent Migration NOTES ↗

Kiro is an AI-powered IDE by AWS that uses "specs" to translate requirements into structured plans, then automatically generates and implements code.

Kiro V3 adds `/upgrade-agent` command to migrate V2 custom agent configs to the universal format

└──▷ **TRY IT**

Convert existing V2 custom agent configs in place so they work in both V2 and V3 sessions.

```
$ /upgrade-agent
```

- New `/upgrade-agent` command migrates V2 custom agent configuration files to the universal format compatible with both V2 and V3 modes.
- Automatic retries for interrupted model responses, reducing manual intervention during long agent sessions.

WAS THIS USEFUL? 👍 👎

Pi ⓘ v0.82.0 NOTES ↗

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

Pi v0.82.0 adds constrained tool sampling, OAuth login for OpenRouter/Kimi Code, and streaming bash RPC events.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.82.0 https://github.com/earendil-works/pi.git
# already have the repo? check out this version:
$ git checkout v0.82.0
```

└──▷ TRY IT

Authenticate with OpenRouter interactively so Pi can use it without manually setting an API key.

```
$ /login
```

Inspect session and model context from inside a bash tool script — useful for writing tools that behave differently per provider or model.

```
$ echo "Session: $PI_SESSION_ID | Provider: $PI_PROVIDER | Model:
$PI_MODEL | Reasoning: $PI_REASONING_LEVEL"
```

- Adds `Tool.constrainedSampling` supporting strict JSON Schema (`prefer` / `require`) and OpenAI Lark/regex grammar variants across OpenAI, Anthropic, Amazon Bedrock, Google Gemini, and Mistral providers.
- Adds model capability flags `supportsGrammarTools` and `supportsStrictTools` to gate constrained sampling requests to models that support them.
- Adds OAuth PKCE login for OpenRouter via `/login`, automatically minting a user-controlled API key without manual configuration.
- Adds OAuth device-authorization login for Kimi Code subscriptions via `/login`, with automatic token refresh.
- Exposes `PI_SESSION_ID`, `PI_SESSION_FILE`, `PI_PROVIDER`, `PI_MODEL`, and `PI_REASONING_LEVEL` as environment variables inside commands run by built-in and factory-created bash tools.

+1 more

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ rust-v0.146.0-alpha.6

3 RELEASES · 2026-07-23 → 2026-07-24 NOTES ↗ CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI alpha.6 adds remote code-mode host support, nonblocking TUI interrupts, and skill-catalog context warnings.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.146.0-alpha.6
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.146.0-alpha.6
```

└──▷ TRY IT

Connect the app-server to a remote code-mode host over a secure WebSocket instead of the default in-process host.

```
$ codex-app-server --code-mode-host wss://example.com/code-mode --
listen ws://127.0.0.1:9000
```

- Adds `--code-mode-host` flag to the app-server to connect to a remote code-mode host over WebSocket (`wss://`).
- Supports WebSocket transport for the code-mode host in the app-server, enabling remote code-mode sessions.
- Exposes remote skill icon URLs through the app server.
- Exposes Browser Use requirements through the app server.
- Routes exec-server HTTP and WebSocket traffic through the configured proxy policy.

+7 more

▾ 2 more releases in this issue · 2026-07-23 → 2026-07-24

> rust-v0.146.0-alpha.3.1 NOTES ↗ CODE ↗ PRE-RELEASE
>
> Custom model providers can now opt into standalone web search via a new `supports_standalone_web_search` setting.
>
> └──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.146.0-alpha.3.1
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.146.0-alpha.3.1
```

- Adds `supports_standalone_web_search` model-provider setting that lets custom Responses API providers opt into the standalone `web.run` web search tool.
- Routes standalone web search requests through the custom provider's own endpoint and authentication when the setting is enabled.

`rust-v0.146.0-alpha.5` NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI alpha.5 adds thread pinning, MCP tool prefix omission, custom provider web search, and MCP connection reuse.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.146.0-alpha.5
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.146.0-alpha.5
```

- Adds `isPinned` field to thread responses with pin/unpin support via `thread/metadata/update` and an `isPinned` filter on `thread/list`, enabling persistent thread pinning with cursor-based pagination.
- Allows omitting MCP tool name prefixes on a per-server basis, reducing verbosity when working with trusted MCP servers.
- Allows custom providers to opt into standalone web search, extending web-search capability beyond first-party models.
- Reuses MCP connections across runtime refreshes and replaces closed connections during reconciliation, reducing reconnection overhead.
- Allows disabling the multi-agent wait tool, giving operators control over multi-agent coordination primitives.

+4 more

WAS THIS USEFUL? 👍 👎

Cline ⓘ desktop-v0.0.4 NOTES ↗ CODE ↗

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

Cline desktop v0.0.4 adds workspace-free chat, file drag-and-drop, one-time routines, and agentic compaction by default.

└──▷ GET THIS VERSION

```
$ git clone --branch desktop-v0.0.4
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout desktop-v0.0.4
```

└──▷ TRY IT

Create a headless scheduled routine that runs unattended in YOLO mode — no human approval prompts.

```
$ cline schedule create "Nightly dependency audit" --cron "0 2 * * *"
--prompt "Check for outdated dependencies and open a PR if any are
found" --workspace /path/to/repo --mode yolo
```

- Supports chat sessions without opening a project folder (workspace-free mode).
- Drag and drop files directly onto the chat to attach them.
- Image attachments display inline in the chat transcript.
- Schedule one-time (non-recurring) routines with navigation to jump to a routine's run.
- Agent sessions now use agentic compaction by default to keep long conversations within context.

+6 more

WAS THIS USEFUL? 👍 👎

Goose ⓘ v1.44.0 NOTES ↗ CODE ↗

an open source, extensible AI agent that goes beyond code suggestions - install, execute, edit, and test with any LLM

Goose v1.44.0 adds new model support (GPT-5.6, Gemini, Muse Spark 1.1), new providers, `--edit` session flag, per-message usage stats, and OllamaCloud integration.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.44.0 https://github.com/aaif-goose/goose.git
# already have the repo? check out this version:
$ git checkout v1.44.0
```

└──▷ TRY IT

Edit a prior session's conversation before forking it — useful for pruning context or replaying with adjusted prompts.

```
$ goose session --edit
```

- Adds support for latest Gemini models.
- Adds Muse Spark 1.1 support via the Meta Models API.
- Adds MiniMax-M3 and missing M2.7 model variants.
- Adds declarative Sakana AI provider for the OpenAI-compatible Fugu API.
- Adds iFlytek Spark and Astron MaaS providers.

+10 more

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

Ollama ⓘ v0.32.3 NOTES ↗ CODE ↗

Get up and running with Kimi-K2.6, GLM-5.2, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.

Ollama v0.32.3 adds CUDA on Windows ARM64, B200/CUDA 12 support, Laguna 2.1 model capabilities, and lower iGPU memory use.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.32.3 https://github.com/ollama/ollama.git
# already have the repo? check out this version:
$ git checkout v0.32.3
```

- Adds CUDA support on Windows ARM64 devices.
- Adds NVIDIA B200 GPU support via CUDA 12.
- Reduces memory usage on Linux CUDA and ROCm integrated GPUs.
- Adds chat, thinking, and tool calling support for Laguna 2.1 models.
- Restores Claude Code Channels support when launching via `ollama launch claude`.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 24, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe