THE CYBER TOOLCHAIN NO. 013

# Head

THE WEEK'S TOP PICKS

PUBLISHED SEPTEMBER 25, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick



*The best tooling updates that shipped this week.*

▸ // HOW THESE PICKS ARE MADE SHOW

▾ // THE BRIEF HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 GPT-5.5

What stands out in this week's picks, grouped by what it lets you do. Every tool named links to its pick below.

Three unrelated things stand out today: Semgrep can notify and rescan org-wide when malicious packages appear, Microsoft Entra can block sensitive uploads to unsanctioned AI tools at the network layer, and NETRESEC can classify possible C2 in PCAPs without byte signatures.

PROTECT

**Stop malicious dependencies before or after they land**

Semgrep adds org-wide policy, rescans, and early blast-radius notifications for zero-day package malware, so teams can find affected projects without waiting for the next normal scan. Cycode moves package blocking onto developer workstations, stopping risky installs before they become pipeline or repository cleanup work.

Semgrep · Cycode

PROTECT

**Enforce AI data-use policy where uploads leave the network**

Microsoft Entra Global Secure Access now ties Purview classification and DLP decisions to network-layer access, so sensitive uploads to unsanctioned AI tools can be stopped even when the application itself is outside the managed SaaS stack.

Microsoft Entra

DETECT

**Classify suspicious PCAP traffic without waiting for signatures**

NETRESEC's FlowCarp uses statistical application-layer classification and a `flowcarp` CLI for custom models, giving analysts a way to flag C2-like protocols when there is no byte-sequence signature. Netcap's expanded OT/ICS decoding makes DNP3 , Modbus , and Siemens S7comm traffic easier to label and inspect during packet analysis.

NETRESEC · Netcap

GOVERN

**Put AI agents under identity and access oversight**

Okta added discovery, gateway enforcement, scoped agent-to-agent access, SSO , revocation, and certification for AI agents. That gives identity teams a control plane for agent access instead of treating agents as unmanaged scripts with long-lived permissions.

Okta Identity Engine

Does **Opus 5** read better? 👍 👎

VIEW

ISSUE VIEW

**scroll**



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

**everything**



BRIEFED | **EVERYTHING**

Every change in this week's picks, in full.

**all 42 changes**

Do you prefer **scroll**? 👍 👎

// THE PICKS

| | | |
|---|---|---|
| 01 | **SonarQube** | SAST / SCA / Code Security |
| 02 | **Palo Alto Networks** | Firewalls & Network Defense |
| 03 | **Semgrep** | SAST / SCA / Code Security |
| 04 | **Docker Sandboxes** | AI/LLM Security |
| 05 | **Cycode** | SAST / SCA / Code Security |

# 01 SonarQube

announced · Sep 21, 2026

SAST / SCA / Code Security

Trusted by 7M+ developers, SonarQube is the AI code review and code verification platform that catches quality and security issues before they reach production.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

SonarQube Hunter Agent adds AI-driven detection of business-logic, broken access control, and authentication/session flaws that static analysis and SCA usually miss. It lands findings as normal SonarQube issues with severity, file, line, triage, scheduling, and an auditable evidence trail, making it immediately usable in existing AppSec workflows.

> SonarQube Cloud added first-class Rust analysis, including Clippy -based rules, Rust metrics, coverage, duplication detection, quality gates, and scanner properties for Cargo and Clippy control. The window also brought Hunter Agent AI verification, native R analysis, and zero-configuration GitLab automatic analysis.

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

01 **Rust analysis in SonarQube Cloud** NEW 96

SonarQube Cloud now supports Rust as a first-class language with 122 curated Clippy rules enforced through quality profiles, plus library- and framework-aware Rust rules for tokio async handling, serde serialization, and actix-web misuses. It calculates Cognitive Complexity and Cyclomatic Complexity, imports LCOV or Cobertura coverage into the issue view, detects duplicate Rust code with CPD, analyzes Rust in the same multi-language project view as C/C++, Java, and JavaScript, and supports quality gate enforcement for Rust with pass/fail thresholds scoped to new code for CI blocking. The scanner adds `sonar.rust.cargo.manifestPaths` for non-default `Cargo.toml` locations and `sonar.rust.clippy.enabled`, which can be set to `false` to import an external Clippy JSON report instead of running `cargo clippy` automatically.

Analyze a Rust workspace where Cargo.toml is not at the repo root — point the scanner at the correct manifest path.

```
$ sonar-scanner -
Dsonar.rust.cargo.manifestPaths=services/myservice/Cargo.toml
```

Use a pre-generated Clippy JSON report (e.g. from a prior CI step) instead of letting SonarQube run cargo clippy itself.

```
$ sonar-scanner -Dsonar.rust.clippy.enabled=false
```

— Exact scanner properties and runnable examples `launch-20260922-39805485`

02 **Hunter Agent AI verification** NEW 74

Hunter Agent adds an AI-driven verification layer that detects business-logic, broken access control, and authentication/session vulnerabilities through structured multi-stage playbooks. Findings appear as standard SonarQube issues with severity, file, and line in the existing triage, assignment, and tracking workflow, are summarized on a dedicated Hunter Agent page, can run manually or daily against long-lived branches, and produce an auditable evidence trail mapped to DORA, NIS2, PCI DSS 4.0.1, HIPAA, GDPR, EU CRA, NYDFS, and FDA 21 CFR Part 11 requirements.

— Strong behavior detail, limited operation path `launch-20260921-74341ed0`

03 **GitLab automatic analysis** NEW 67

SonarQube Cloud extends zero-configuration automatic analysis, previously available for GitHub and Azure DevOps, to GitLab. It scans eligible GitLab repositories without `.gitlab-ci.yml` edits, covers 20+ languages including C#, Java, and Python, and starts analysis after users link GitLab, import repositories, and select eligible repositories.

Onboard a GitLab repository to SonarQube Cloud for instant, pipeline-free scanning on every merge request.

> 📍 1. Sign in to SonarQube Cloud and link your GitLab account (ensure your access token has repository and metadata read permissions). 2. Go to 'Import repositories' and select the GitLab repositories you want to analyze. 3. For eligible repositories, automatic analysis starts immediately — no further configuration needed. 4. Results appear on the default branch and active merge requests shortly after import.

— Clear onboarding steps, limited internals `launch-20260915-5d8d3f1f`

04 **Native R language analysis** NEW 65

SonarQube adds native R support on Cloud and Server, analyzing R source files and embedded R code inside R Markdown documents with 82 rules covering key sections of the lintr rule set plus Sonar-developed rules. It imports lintr problem reports, adds Cobertura test coverage import, extends secrets detection, code metrics, syntax highlighting, and copy-paste duplication analysis to R, and applies existing quality profiles, quality gates, PR decoration, and branch analysis to R codebases.

— Names formats and scope, no command `launch-20260915-f7587a71`

announced · Sep 21, 2026

SonarQube Hunter Agent adds AI-driven detection of business logic, access control, and authentication flaws as standard issues.

- Adds Hunter Agent, a new AI-driven verification layer that detects business-logic, broken access control, and authentication/session vulnerabilities through structured multi-stage playbooks — classes of flaw that static analysis and SCA miss.
- Surfaces Hunter Agent findings as standard SonarQube issues (with severity, file, and line) in the existing triage, assignment, and tracking workflow, with results also summarized on a dedicated Hunter Agent page.
- Supports running Hunter Agent manually or on a daily schedule against long-lived branches to suit different team cadences.
- Produces a reproducible, auditable evidence trail from detection through triage and remediation, mapped to DORA, NIS2, PCI DSS 4.0.1, HIPAA, GDPR, EU CRA, NYDFS, and FDA 21 CFR Part 11 requirements.

announced · Sep 22, 2026

SonarQube Cloud adds first-class Rust support: 122 Clippy rules, complexity metrics, coverage import, and quality gates.

**└──▷ TRY IT**

Analyze a Rust workspace where Cargo.toml is not at the repo root — point the scanner at the correct manifest path.

```
$ sonar-scanner -
Dsonar.rust.cargo.manifestPaths=services/myservice/Cargo.toml
```

Use a pre-generated Clippy JSON report (e.g. from a prior CI step) instead of letting SonarQube run cargo clippy itself.

```
$ sonar-scanner -Dsonar.rust.clippy.enabled=false
```

- Adds `sonar.rust.cargo.manifestPaths` config property to point the Rust analyzer at non-default `Cargo.toml` locations.
- Adds `sonar.rust.clippy.enabled` config property (set to `false` to import an external Clippy JSON report instead of running `cargo clippy` automatically).
- Supports Rust as a first-class language in SonarQube Cloud, with 122 curated Clippy rules enforced through quality profiles across all projects.

- Adds library- and framework-aware Rust rules covering tokio async handling, serde serialization, and actix-web misuses — beyond what Clippy detects.
- Adds Cognitive Complexity and Cyclomatic Complexity metric calculation for Rust code.

+4 more

announced · Sep 15, 2026

SonarQube adds native R language support with 82 rules, secrets detection, lintr import, and Cobertura coverage on Cloud and Server.

- Analyzes R source files and embedded R code inside R Markdown documents with 82 rules covering key sections of the lintr rule set plus additional Sonar-developed rules.
- Imports lintr problem reports directly into SonarQube, surfacing existing findings in the centralized view alongside other languages.
- Adds Cobertura test coverage import for R projects.
- Extends secrets detection, code metrics, syntax highlighting, and copy-paste duplication analysis to R.
- Applies existing quality profiles, quality gates, PR decoration, and branch analysis to R codebases.

announced · Sep 15, 2026

SonarQube Cloud adds zero-configuration automatic analysis for GitLab, scanning every merge request and main-branch commit without .gitlab-ci.yml edits.

**└──▷ HOW TO FIND IT**

Onboard a GitLab repository to SonarQube Cloud for instant, pipeline-free scanning on every merge request.

> 📍 1. Sign in to SonarQube Cloud and link your GitLab account (ensure your access token has repository and metadata read permissions). 2. Go to 'Import repositories' and select the GitLab repositories you want to analyze. 3. For eligible repositories, automatic analysis starts immediately — no further configuration needed. 4. Results appear on the default branch and active merge requests shortly after import.

- Extends zero-configuration automatic analysis — previously exclusive to GitHub and Azure DevOps — to GitLab, covering 20+ languages including C#, Java, and Python.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

## 02 Palo Alto Networks

changelog · Sep 22, 2026

Firewalls & Network Defense

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

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Unit 42 Continuous Frontier AI Defense adds always-on agentic offensive testing that continuously finds, validates, and remediates exposures across web apps, APIs, cloud infrastructure, source repositories, and network assets. The Advanced Adversary Simulation validates end-to-end attack paths, closing a practical gap between asset scanning and proven exploitability.

> Palo Alto Networks launched Unit 42 Continuous Frontier AI Defense, an always-on agentic offensive cyber service for continuously finding, validating, and remediating enterprise exposures. The service includes continuous estate testing, multi-model AI routing, and adversary simulation across web apps, APIs, cloud infrastructure, source code repositories, and network assets.

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

### 01 Multi-Model AI Harness NEW

62

The proprietary Multi-Model AI Harness routes offensive security tasks to the model best suited for each job. Palo Alto Networks says this addresses findings that no single model catches more than 40% of vulnerabilities and that leading models share less than 10% overlap in identified exposures.

— Mechanism and metrics, no operator controls 2026-09-22-d8eee514

### 02 Continuous Frontier AI Defense NEW

60

Unit 42 Continuous Frontier AI Defense is a new always-on agentic offensive security service that continuously finds, validates, and remediates enterprise exposures. It uses proprietary multi-model AI harnesses with Anthropic Claude Mythos 5 and OpenAI GPT-5.6-Cyber.

— Names service and models, but no usage surface 2026-09-22-d8eee514

THINNER COVERAGE BELOW

### 03 Continuous Testing Engine NEW

59

The Continuous Testing Engine performs a full-estate baseline scan followed by always-on testing as an organization's environment changes. Coverage includes web apps, APIs, cloud infrastructure, source code repositories, and network assets.

— Clear scope, no access path 2026-09-22-d8eee514

04 **Advanced Adversary Simulation** NEW 53

Advanced Adversary Simulation validates end-to-end attack paths to prove real-world exploitability. It works across first- and third-party web apps, APIs, cloud infrastructure, source code repositories, and network assets.

— Explains validation scope, not how to run 2026-09-22-d8eee514

changelog · Sep 22, 2026

Palo Alto Networks launches Unit 42 Continuous Frontier AI Defense, an always-on agentic offensive security service powered by Anthropic Claude Mythos 5 and OpenAI GPT-5.6-Cyber.

- Introduces Unit 42 Continuous Frontier AI Defense, a new always-on agentic offensive security service that continuously finds, validates, and remediates enterprise exposures using proprietary multi-model AI harnesses with Anthropic Claude Mythos 5 and OpenAI GPT-5.6-Cyber.
- Adds a Continuous Testing Engine that performs a full-estate baseline scan followed by always-on testing as an organization's environment changes, covering web apps, APIs, cloud infrastructure, source code repositories, and network assets.
- Introduces a proprietary Multi-Model AI Harness that routes offensive security tasks to the model best suited for each job, addressing the finding that no single model catches more than 40% of vulnerabilities and that leading models share less than 10% overlap in identified exposures.
- Delivers Advanced Adversary Simulation that proves real-world exploitability by validating end-to-end attack paths across first- and third-party web apps, APIs, cloud infrastructure, source code repositories, and network assets.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

## 03 Semgrep

announced · Sep 24, 2026

SAST / SCA / Code Security

Semgrep is a static analysis tool that scans code for bugs, security vulnerabilities, and policy violations using customizable rules.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Semgrep Supply Chain adds zero-day malware detection that rescans customer repositories within minutes of a supply-chain incident and reports affected repos, projects, findings, and scan progress. The new org-wide response policies and Malware Firewall extend that from detection into automated alerting and pre-install blocking.

> Semgrep Supply Chain added zero-day malware detection and response automation, including org-wide policies, incident rescans, early blast-radius notifications, and a private-beta Malware Firewall. Semgrep also added CSV export for Agentic Workflows issues from the Issues page and public API.

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

01 **Malware Firewall private beta** NEW 75

Semgrep added a Malware Firewall in private beta, embedded in Semgrep Guardian via Agent admin. It intercepts dependency requests at the network level between the package manager and public registries and blocks malicious packages before they reach developer machines, with no per-device install.

— Mechanism and placement are described launch-20260924-77b2b3de

02 **Supply Chain malware response automation** NEW 65

Semgrep Supply Chain now supports org-wide policies that trigger automated responses when a malware finding is generated, including opening a Jira ticket, sending a Slack alert, or calling a webhook with no human in the loop.

— Names response targets, but no setup path launch-20260924-77b2b3de

03 **Supply Chain incident notifications** NEW 64

Customers can opt in to early notifications that proactively report their blast radius via Slack or webhooks when a new supply chain incident is detected. Notifications cover affected repos, projects, findings, and scan progress.

— Clear channels and contents, no configuration details launch-20260924-77b2b3de

04 **Zero-day incident detection** NEW 63

Semgrep added zero-day incident detection that rescans all customer repositories with AI-assisted, expert-verified advisory rules within minutes of a supply chain incident. Findings are posted to the Semgrep dashboard, organized by advisories and rolled up into incidents.

— Explains behavior, but no user control `launch-20260924-77b2b3de`

THINNER COVERAGE BELOW

05 **Agentic Workflows CSV export** NEW 49

Semgrep now supports exporting Agentic Workflows issues as CSV from the Issues page and through the public API.

— Names format and entry points only `snapshot-20260923`

`announced · Sep 24, 2026`

Semgrep Supply Chain adds zero-day malware detection with automated org-wide response policies and a malware firewall (private beta).

- Adds org-wide Supply Chain policies that trigger automated responses — open a Jira ticket, send a Slack alert, or call a webhook — the moment a malware finding is generated, with no human in the loop.
- Adds zero-day incident detection that rescans all customer repositories with AI-assisted, expert-verified advisory rules within minutes of a supply chain incident and posts findings to the Semgrep dashboard organized by advisories rolled up into incidents.
- Adds early-notification opt-in that proactively informs customers via Slack or webhooks of their blast radius — affected repos, projects, findings, and scan progress — when a new supply chain incident is detected.
- Adds a Malware Firewall (private beta), embedded in Semgrep Guardian via Agent admin, that intercepts dependency requests at the network level between the package manager and public registries and blocks malicious packages before they reach developer machines, with no per-device install.

`snapshot-20260923`

Semgrep adds CSV export for Agentic Workflows issues via the Issues page and public API.

- Supports exporting Agentic Workflows issues as CSV from the Issues page and through the public API.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

# 04 Docker Sandboxes

announced · Sep 24, 2026

AI/LLM Security

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Docker Sandboxes adds Cloud Sandboxes and an OCI-native Kit format that lets teams run coding agents in managed microVMs with declared network and credential authority. Kits, proxy-managed secrets, network policies, and authority-diff gating make agent execution governable rather than just isolated.

> Docker Sandboxes launched managed Cloud Sandboxes and the v3 OCI Kit format for packaging sandbox authority. The window also added model-provider configuration, MCP transport controls, network-policy hardening, richer CLI/JSON surfaces, and several breaking CLI behavior changes.

└──▷ **WHAT SHIPPED · 26 FEATURES** most completely described first ? what's the number?

01 **OCI v3 Kit specification** NEW 96

Docker Sandbox Kit Specification v3 defines a portable OCI image format, open source under Apache 2.0 at `docker/sandbox-kit-spec`, for declaring what authority a sandbox agent may exercise. Kit declarations live in the OCI manifest annotation `vnd.docker.sandbox.kit.descriptor`, build with `docker buildx build`, pull with `docker pull`, and are pinnable by digest; v3 kits separate `workload` Kits, which supply the root filesystem, from `mixin` Kits, which overlay a CLI, network rule, or credential binding. The spec defines versioned capability types including `com.docker.sandbox/network-policy@2` and `com.docker.sandbox/credential@1`, plus conformance suites for artifacts and runtimes.

Run a sandboxed agent with the GitHub CLI mixin to give it scoped GitHub access defined by the Kit's network-policy and credential declarations.

```
$ cd examples
sbx run ./hello --kit ./gh
```

Author a mixin Kit that allows most GitHub API methods but explicitly denies DELETE on `/repos`/**, keeping a token that can open PRs from deleting repositories.

```yaml
capabilities:
  - type: com.docker.sandbox/network-policy@2
    config:
      runtime:
        allow:
          - github.com
          - hosts: [api.github.com]
            methods: [GET, HEAD, POST, PATCH, PUT, DELETE]
        deny:
          - hosts: [api.github.com]
            methods: [DELETE]
            paths: [/repos/**]
  - type: com.docker.sandbox/credential@1
    optional: true
    config:
      service: github
      phase: runtime
      apiKey:
        name: GH_TOKEN
      proxyManaged: true
      inject:
        - {domain: api.github.com, header: Authorization,
format: 'Bearer %s'}
```

— Exact spec surfaces and build commands `v0.45.0` `launch-20260924-0c0143b0`

---

02 **Cloud Sandboxes compute** NEW 94

Cloud Sandboxes run sandboxes on Docker-managed cloud compute with `sbx --cloud run <kit>` and support moving local work with `sbx move --to cloud`. Size options range from Micro (1 vCPU / 2 GiB, $0.07/hr) to XL (16 vCPU / 32 GiB, $1.12/hr), sessions can run up to 24 hours, and cloud support requires `sbx` version 0.45.1 or later.

Launch a Claude Code agent in the cloud before leaving for the day so it keeps running overnight without your laptop.

```
$ sbx --cloud run claude
```

Hand off a long-running refactor you started locally to cloud compute so you can close your laptop without interrupting the agent.

```
$ sbx move my-project --to cloud
```

— Commands, limits, pricing, version requirement `launch-20260924-6c1aef17`

03 **Kit authority enforcement** NEW 93

V3 Kit capabilities can scope network access by HTTP method and path, declare install-phase network access and credential use, and specify where an agent reads shared skills. Network-policy capabilities support explicit `allow` and `deny` rules scoped to `hosts`, `methods`, and `paths`, with `deny` taking precedence, while credential capabilities support `proxyManaged: true`, injecting real tokens into requests to named domains and exposing only a sentinel inside the sandbox. Authority-diff gating reduces every descriptor to a normalized grant set so a runtime may auto-apply updates that stay within the recorded set, but any widening, including removing a `deny` rule, requires explicit approval.

Author a mixin Kit that allows most GitHub API methods but explicitly denies DELETE on `/repos/**`, keeping a token that can open PRs from deleting repositories.

yaml

```yaml
capabilities:
  - type: com.docker.sandbox/network-policy@2
    config:
      runtime:
        allow:
          - github.com
          - hosts: [api.github.com]
            methods: [GET, HEAD, POST, PATCH, PUT, DELETE]
        deny:
          - hosts: [api.github.com]
            methods: [DELETE]
            paths: [/repos/**]
  - type: com.docker.sandbox/credential@1
    optional: true
    config:
      service: github
      phase: runtime
      apiKey:
        name: GH_TOKEN
      proxyManaged: true
      inject:
        - {domain: api.github.com, header: Authorization,
format: 'Bearer %s'}
```

— Concrete capability fields and enforcement behavior `v0.45.0` `launch-20260924-0c0143b0`

04 **Network policy protocol selection** NEW 88

Local network allow rules gained `--protocol`, allowing UDP to be selected explicitly alongside the default TCP coverage. Experimental outbound UDP now follows sandbox network policy.

Allow outbound UDP (e.g. for a DNS-over-UDP workflow) in a sandbox network policy instead of the TCP-only default.

```
$ sbx policy allow network --protocol udp 8.8.8.8:53
```

— Exact flag and behavior `v0.45.0`

05 **Network policy bypass fixes** IMPROVED 87

Reverse-DNS lookups are now restricted to destination IPs already authorized by policy, closing a previous policy bypass. DNS resolution is blocked when no network rule permits it, and `sbx policy allow network`, `sbx policy deny network`, `--allow-network`, and `--deny-network` now reject malformed patterns before saving them.

— Commands and before-after policy behavior `v0.45.0`

06 **Non-interactive destructive confirmations** BREAKING 86

Destructive commands gained `--force` for non-interactive confirmation, and `sbx kit builder history rm` also gained `--yes` and `-y`. Commands that remove resources now prompt for confirmation and return a non-zero exit code when the prompt is declined.

— Flags and changed exit behavior `v0.45.0`

07 **MCP catalog removal** BREAKING 86

`sbx mcp catalog` has been removed. Register a remote MCP server with `sbx mcp add` and then authorize it with `sbx mcp auth` instead.

— Removed command and replacement commands `v0.45.0`

08 **Custom model providers** NEW 85

`sbx run --model` now supports any OpenAI- or Anthropic-compatible endpoint configured through the new `model.providers` setting. For `sbx run opencode --model`, the model's supported thinking levels are exposed as OpenCode variants selectable with Ctrl+T.

Route agent runs through a custom Anthropic-compatible endpoint instead of the built-in provider.

```
$ sbx run --model my-anthropic-proxy/claude-3-5-sonnet
'summarize the repo'
```

— Flag, setting, compatible endpoints `v0.45.0`

09 **Remote MCP HTTP/2 control** NEW 84

`sbx mcp add --disable-http2` disables HTTP/2 for a remote MCP transport. The flag is intended for servers whose HTTP/2 handling stalls long-lived streams.

Work around a remote MCP server that stalls long-lived HTTP/2 streams by forcing HTTP/1.1.

```
$ sbx mcp add --disable-http2 my-mcp-server
https://mcp.example.com/sse
```

— Exact command flag and use case `v0.45.0`

10 **Missing resource errors** BREAKING 82

`sbx secret rm` now returns an error on stderr when the requested secret does not exist, where it was previously silent. `sbx mcp rm` now returns an error when the requested MCP server is not registered, where it was previously silent.

— Exact commands and changed failure mode `v0.45.0`

11 **Kit validation hardening** IMPROVED 80

`sbx kit validate` now rejects malformed API-key declarations, including invalid names, missing injection domains, invalid format placeholders, and Basic-auth usernames containing a colon. It also warns for no-op or out-of-allowlist declarations.

— Validator command and rejection cases `v0.45.0`

12 **Detached exec failure** BREAKING 79

`sbx exec -d` and `--detach` now fail immediately for unsupported detached execution instead of silently running in the foreground.

— Exact flags and before-after behavior `v0.45.0`

13 **Kit sets and composition** NEW 77

Kit sets let authors combine a workload and mixins, pin component versions, and publish the result as a single OCI reference with its own capabilities, lifecycle hooks, instructions, and arguments. A `kind: set` descriptor names a collection of Kits, runs coherence checks at build time, and merges them into one ordinary Kit so incoherence is caught at the author's build rather than at launch.

— Names descriptor field and behavior `v0.45.0` `launch-20260924-0c0143b0`

14 **Sandbox JSON metadata** IMPROVED 72

`sbx ls --json` and `sbx inspect --json` now include `created_at`, CPU limits, and memory limits in their output.

— Exact commands and JSON field `v0.45.0`

15 **Cloud Kit catalog** NEW 71

Cloud Sandboxes introduced pre-configured, pre-built Kits for leading coding agents, launchable with `sbx --cloud run <kit>`. Named Kits include Claude Code, Codex, Copilot, Antigravity, Open Code, and Hermes.

Launch a Claude Code agent in the cloud before leaving for the day so it keeps running overnight without your laptop.

```
$ sbx --cloud run claude
```

— Command and supported kit names `launch-20260924-6c1aef17`

16 **Policy listing creation filter** NEW 67

`sbx policy ls` gained the `--created-via` filter. The command also now shows how each rule was created.

— Exact command and filter `v0.45.0`

17 **Name-only list output** NEW 65

MCP server, secret, skill, template, volume, and policy-profile list commands gained `--quiet` and `-q` for name-only output.

— Exact flags, broad command family `v0.45.0`

18 **Signed macOS app bundle** IMPROVED 60

macOS distributions now ship as a single signed `Sbx.app` bundle. Homebrew and tarball PATH installs work through a symlink into the bundle.

— Package name and install behavior `v0.45.0`

THINNER COVERAGE BELOW

19 **Cloud secrets proxy** NEW 59

Cloud Sandboxes added a secrets proxy that injects stored keys or tokens per request so agents do not hold the actual secret value. Separately, local and cloud sandboxes maintain separate secrets, templates, and network policies, so existing local workflows must be reconfigured before moving to cloud.

— Mechanism plus migration caveat `launch-20260924-6c1aef17`

20 **Centralized MCP gateway** NEW 58

A centralized MCP gateway lets agents running in cloud, locally, or in other clients such as ChatGPT desktop share one set of connected MCP servers. The launch names Jira, Linear, Grafana, incident.io, and any streamable HTTP endpoint as supported server targets.

— Named integrations, no setup surface `launch-20260924-6c1aef17`

21 **OAuth credential coordination** IMPROVED 56

Multiple OAuth-backed agents can now be composed in the same sandbox with credentials scoped to the individual kits that request them. OAuth refreshes are coordinated across sandboxes that share credentials, preventing simultaneous refreshes from forcing another sign-in.

— Behavior described, no user surface `v0.45.0`

22 **V3 kit compatibility break** BREAKING 56

V3 workloads and mixins must be used together. They cannot be combined with v1 or v2 kits.

— Compatibility constraint only `v0.45.0`

23 **OpenCode GitHub Copilot setup** IMPROVED 52

The OpenCode kit now configures GitHub Copilot from the account's stored GitHub credential, enabling Copilot models without a separate device login.

— Integration named, no command `v0.45.0`

24 **Cloud network policies** NEW 49

Cloud Sandboxes added network policies that define which endpoints agents can reach. Policies are configured once and applied across all cloud sandboxes, but local and cloud sandboxes maintain separate network policies and require reconfiguration before moving workflows to cloud.

— Scope described, no config syntax `launch-20260924-6c1aef17`

25 **Web console for Cloud Sandboxes** NEW 45

A web console can spin up and manage Cloud Sandboxes without using the CLI. The console includes Kit selection and a Run button.

— UI starting point only `launch-20260924-6c1aef17`

26 **Global service secret updates** IMPROVED 45

Adding, updating, or removing global service secrets now updates existing local sandboxes without a restart.

— Clear behavior, no command named `v0.45.0`

└──▷ **BREAKING ON UPGRADE**

! Local and cloud sandboxes maintain separate secrets, templates, and network policies — existing local workflows must be reconfigured before moving to cloud.

! `sbx mcp catalog` has been removed; register a remote MCP server with `sbx mcp add` and then authorize it with `sbx mcp auth` instead.

! `sbx secret rm` now returns an error on stderr when the requested secret does not exist (previously silent).

! `sbx mcp rm` now returns an error when the requested MCP server is not registered (previously silent).

! V3 workloads and mixins must be used together; they cannot be combined with v1 or v2 kits.

! Commands that remove resources now prompt for confirmation and return a non-zero exit code when the prompt is declined; use `--force` or `--yes` / `-y` in non-interactive workflows.

! `sbx exec -d` / `--detach` now fails immediately for unsupported detached execution instead of silently running in the foreground.

`announced · Sep 24, 2026`

Docker Sandboxes adds Cloud Sandboxes: managed microVM compute in the cloud with `sbx --cloud run` and `sbx move --to cloud`.

.
.
.
.

└──▷ **TRY IT**

Launch a Claude Code agent in the cloud before leaving for the day so it keeps running overnight without your laptop.

```
$ sbx --cloud run claude
```

Hand off a long-running refactor you started locally to cloud compute so you can close your laptop without interrupting the agent.

```
$ sbx move my-project --to cloud
```

› Adds `sbx --cloud run <kit>` to launch a sandbox on Docker-managed cloud compute, with size options from Micro (1 vCPU / 2 GiB, $0.07/hr) to XL (16 vCPU / 32 GiB, $1.12/hr) and sessions up to 24 hours.

› Introduces Kits: pre-configured, pre-built sandbox images for leading coding agents (Claude Code, Codex, Copilot, Antigravity, Open Code, Hermes), launchable with a single `sbx --cloud run <kit>` command.

› Adds a centralized MCP gateway so agents running in cloud, locally, or in other clients (e.g. ChatGPT desktop) can share a single set of connected MCP servers (Jira, Linear,

Grafana, incident.io, or any streamable HTTP endpoint).

› Adds a Cloud Sandboxes secrets proxy that injects stored keys or tokens per-request so agents never hold the actual secret value.

› Adds network policies that define which endpoints agents can reach, configurable once and applied across all cloud sandboxes.

+2 more

**└──▷ BREAKING ON UPGRADE**

! Local and cloud sandboxes maintain separate secrets, templates, and network policies — existing local workflows must be reconfigured before moving to cloud.

announced · Sep 24, 2026

Docker Sandbox Kit Specification v3 introduces OCI -native Kit format for declaring agent authority as versioned, reviewable config

**└──▷ TRY IT**

Run a sandboxed agent with the GitHub CLI mixin to give it scoped GitHub access defined by the Kit's network-policy and credential declarations.

```
$ cd examples
sbx run ./hello --kit ./gh
```

Author a mixin Kit that allows most GitHub API methods but explicitly denies DELETE on `/repos`/**, keeping a token that can open PRs from deleting repositories.

```yaml
capabilities:
  - type: com.docker.sandbox/network-policy@2
    config:
      runtime:
        allow:
          - github.com
          - hosts: [api.github.com]
            methods: [GET, HEAD, POST, PATCH, PUT, DELETE]
        deny:
          - hosts: [api.github.com]
            methods: [DELETE]
            paths: [/repos/**]
  - type: com.docker.sandbox/credential@1
    optional: true
    config:
      service: github
      phase: runtime
      apiKey:
        name: GH_TOKEN
      proxyManaged: true
      inject:
        - {domain: api.github.com, header: Authorization, format:
'Bearer %s'}
```

- Introduces the Docker Sandbox Kit Specification v3, open source under Apache 2.0 at `docker/sandbox-kit-spec`, defining a portable OCI image format for declaring exactly what authority a sandbox agent may exercise.
- Stores all Kit declarations in a single OCI manifest annotation (`vnd.docker.sandbox.kit.descriptor`), so Kits build with `docker buildx build`, pull with `docker pull`, and are pinnable by digest.
- Supports two Kit kinds: a `workload` Kit (supplies the root filesystem) and a `mixin` Kit (an overlay adding a CLI, network rule, or credential binding), composed at launch with one workload and any number of mixins.
- Defines versioned capability types — including `com.docker.sandbox/network-policy@2` and `com.docker.sandbox/credential@1` — each with a normative page describing what a conforming runtime must enforce.
- Network-policy capabilities support explicit `allow` and `deny` rules scoped to `hosts`, `methods`, and `paths`, with `deny` taking precedence, so a mixin can permit most of an API while blocking specific destructive operations (e.g. `DELETE /repos/**`).

+4 more

v0.45.0

Docker Sandboxes v0.45.0 ships v3 OCI kits with workload/mixin composition, new model provider settings, and a raft of new CLI flags.

### └──▷ GET THIS VERSION

```
$ git clone --branch v0.45.0 https://github.com/docker/sbx-releases.git
# already have the repo? check out this version:
$ git checkout v0.45.0
```

### └──▷ TRY IT

Route agent runs through a custom Anthropic-compatible endpoint instead of the built-in provider.

```
$ sbx run --model my-anthropic-proxy/claude-3-5-sonnet 'summarize the repo'
```

Allow outbound UDP (e.g. for a DNS-over-UDP workflow) in a sandbox network policy instead of the TCP-only default.

```
$ sbx policy allow network --protocol udp 8.8.8.8:53
```

Work around a remote MCP server that stalls long-lived HTTP/2 streams by forcing HTTP/1.1.

```
$ sbx mcp add --disable-http2 my-mcp-server https://mcp.example.com/sse
```

- Adds `sbx run --model` support for any OpenAI- or Anthropic-compatible endpoint configured via the new `model.providers` setting.
- Adds `sbx mcp add --disable-http2` to disable HTTP/2 for a remote MCP transport, working around servers whose HTTP/2 handling stalls long-lived streams.
- Adds `--quiet` (`-q`) flag to MCP server, secret, skill, template, volume, and policy-profile list commands for name-only output.
- Adds `--protocol` flag to local network allow rules so UDP can be selected explicitly alongside the default TCP coverage.
- Adds `--created-via` filter to `sbx policy ls`, which also now shows how each rule was created.

+16 more

**└──▷ BREAKING ON UPGRADE**

- ! `sbx mcp catalog` has been removed; register a remote MCP server with `sbx mcp add` and then authorize it with `sbx mcp auth` instead.
- ! `sbx secret rm` now returns an error on stderr when the requested secret does not exist (previously silent).
- ! `sbx mcp rm` now returns an error when the requested MCP server is not registered (previously silent).
- ! V3 workloads and mixins must be used together; they cannot be combined with v1 or v2 kits.
- ! Commands that remove resources now prompt for confirmation and return a non-zero exit code when the prompt is declined; use `--force` or `--yes` / `-y` in non-interactive workflows.
- ! `sbx exec -d` / `--detach` now fails immediately for unsupported detached execution instead of silently running in the foreground.

---

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

# 05 Cycode

Introducing Cycode Workstation Protection: ADLC Security Now Stops Software Supply Chain Attacks (2026-09-23)

SAST / SCA / Code Security

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Cycode Workstation Protection intercepts package install commands on developer machines in real time, enforcing cooldown policies and blocking known malicious packages before installation. That closes a supply-chain control gap upstream of CI by applying policy at the workstation regardless of package manager version or local config.

> Cycode added Workstation Protection to stop risky package installs directly on developer workstations before code enters the pipeline. The release adds real-time install interception, cooldown policy enforcement, malicious package intelligence, and MDM deployment without changing developer install workflows.

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

01 **Workstation package install enforcement** NEW 75

Workstation Protection intercepts package install commands in real time on developer workstations to enforce cooldown, or release-age gating, policies and block known malicious packages before installation. It checks every install against a continuously updated threat intelligence feed, including typosquatting, slopsquatting, and persistent compromised maintainer releases, and enforces policies at the device level regardless of package manager version or local config correctness.

– Explains behavior and scope, but no commands or UI path

Introducing Cycode Workstation Protectio…

THINNER COVERAGE BELOW

02 **MDM deployment for workstation protection** NEW 47

Workstation Protection is deployed via Mobile Device Management, with no new console, no new infrastructure, and no change to developer install workflows.

– Names deployment route, but lacks setup steps Introducing Cycode Workstation Protectio…

03 **Workstation signals in Context Intelligence Graph** IMPROVED

Cycode extended ADLC Security coverage upstream to the developer workstation by unifying workstation package enforcement with existing AI Bill of Materials, prompt guardrails, pipeline integrity, and secret scanning into the Context Intelligence Graph.

— Names connected areas, but not user-operable `Introducing Cycode Workstation Protectio…`

> `Introducing Cycode Workstation Protection: ADLC Security Now Stops Software Supply Chain Attacks (2026-09-23)`
>
> Cycode Workstation Protection adds real-time package install enforcement with cooldown policies and malicious package threat intelligence.
>
> - Adds Workstation Protection, a new capability that intercepts package install commands in real time to enforce cooldown (release-age gating) policies and block known malicious packages before they install on developer workstations.
> - Deploys Workstation Protection via Mobile Device Management (MDM) with no new console, no new infrastructure, and no change to developer install workflows.
> - Checks every package install against a continuously updated threat intelligence feed to catch malicious packages that have aged past cooldown windows, including typosquatting, slopsquatting, and persistent compromised maintainer releases.
> - Enforces cooldown policies at the device level, so gating applies regardless of which package manager version a developer or coding agent is running or whether a local config file was configured correctly.
> - Extends ADLC Security coverage upstream to the developer workstation, unifying workstation package enforcement with existing AI Bill of Materials, prompt guardrails, pipeline integrity, and secret scanning into the Context Intelligence Graph.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · September 25, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe