THE CYBER TOOLCHAIN NO. 066

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 6, 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

INFORMATION DENSITY

everything



BRIEFED | **EVERYTHING**

Every tool in the issue, in full. **all 18 changes**

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 4 tools

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

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

Three unrelated things stand out. skill-scanner can now shadow-run a rule change against real findings before it blocks anything, and catches

subprocess calls hidden behind aliases. `safe2 doctor` walks an environment, finds the agent harnesses actually running on it, and reports posture. DOMPurify hardens DOM clobbering in XML content.

PROTECT

Test a policy change on live findings before it blocks a build

skill-scanner 2.1.0 correlates findings through a bounded CEL layer with off/shadow/enforce modes and a `validate-rules` subcommand, so a rule pack can be run against real scan output and inspected before enforcement — the usual way this goes wrong is a new rule failing every pipeline at once. ai-safe2-framework adds matching policy enforcement, schema/evidence and feedback subcommands on the other side of the same problem.

skill-scanner · ai-safe2-framework

IDENTIFY

Find out which agent harnesses are actually running in an environment

`safe2 doctor` discovers agent harnesses, assesses environment posture and emits Decision Cards. Agent tooling tends to arrive per-developer rather than per-platform, so the first honest answer to "what is running here" usually requires someone to go look; this makes it a command.

ai-safe2-framework

PROTECT

Close DOM clobbering paths in sanitized XML

3.4.15 hardens against DOM clobbering in XML content and fixes assorted sanitization edge cases. Clobbering attacks smuggle named elements past a sanitizer to shadow properties the surrounding script trusts, so the payload never has to look like script at all.

DOMPurify

GOVERN

Reopen a submitted assessment instead of starting a new one

v4.0.1 lets teams reopen submitted requirement assessment assignments and adds a URL custom field type. A submitted-is-final workflow means one wrong answer costs a fresh assignment and a broken audit trail; this keeps the correction attached to the original.

CISO Assistant

Does **Opus 5** read better? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Cisco skill-scanner

1 RELEASE · 2026-09-05 NOTES ↗ CODE ↗

RANK

The skill-scanner is a CLI security scanner for AI agent skills.

> skill-scanner 2.1.0 introduces a bounded CEL decision layer for correlating scan findings with off/shadow/enforce modes, adds administrator-trusted and ATR rule pack loading with a new validate-rules subcommand, and expands Python static detection to catch obfuscated and aliased subprocess/file-access patterns.

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

01 CEL-based correlation decision layer NEW

95

skill-scanner adds a bounded CEL decision layer (backed by `cel-go` v0.32.0) that correlates typed analyzer facts after deterministic detection and before optional LLM/Meta analysis, using a versioned protobuf fact contract (`skill_scanner.semantic.v1.ScanFacts`) with strict schema-v2 pack validation, immutable compiled rule generations, expression hashes, and fail-open telemetry. The mode is controlled globally via the `cel.mode` config key (`off` , `shadow` , `enforce`) and the `analyzers.correlation` config key, which enables bounded source/sink and staged-behavior correlation analysis, or per-invocation via the `--cel-mode off|shadow|enforce` (also documented as `--cel-mode MODE`) CLI flag, letting a run be promoted to enforce without changing global configuration.

Promote CEL gates to enforce mode for a single CI run to validate coverage before rolling it out globally.

```
$ skill-scanner scan --cel-mode enforce <target>
```

— Names flags, config keys, library version and fact-contract schema `2.1.0`

02 Trusted and ATR rule pack loading via validate-rules NEW

80

Adds the `--rule-packs atr` flag to opt in to the ATR rule pack as an optional input alongside the core + CEL release gate, and the repeatable `--trusted-rule-pack PATH` flag to load administrator-trusted schema-v2 signature/YARA/CEL packs. The `validate-rules` subcommand is extended to validate both bundled rules and these trusted schema-v2 packs, including CEL parse/type checks, e.g. `skill-scanner validate-rules --trusted-rule-pack <path>` .

Include the ATR rule pack alongside core rules to extend detection coverage in an on-demand scan.

```
$ skill-scanner scan --rule-packs atr <target>
```

Validate an administrator-approved YARA/CEL pack before deploying it to production, checking for schema-v2 compliance and CEL type errors.

```
$ skill-scanner validate-rules --trusted-rule-pack /path/to/org-pack
```

Run in shadow mode to observe CEL policy matches without enforcement — useful when rolling out a new rule pack to catch false positives first.

```
$ skill-scanner validate-rules --trusted-rule-pack /path/to/org-pack --cel-mode shadow
```

Run CEL decisions in shadow mode first to observe results without blocking scans before committing to enforce

```
$ skill-scanner validate-rules --trusted-rule-pack /path/to/org-pack
```

— Flags and subcommand named with runnable example commands `product docs`

THINNER COVERAGE BELOW

03 **Expanded Python static detection** IMPROVED 55

Static detection now identifies shell-enabled subprocess calls made through aliases and class scopes, constructed sensitive-file reads, import aliases, and commands hidden by supported repeating-XOR encoding, plus pipeline detection for recovered runtime commands.

— Describes concrete detection cases but no flag or command to invoke `2.1.0`

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Cure53 DOMPurify ⓘ 1 RELEASE · 2026-09-06 NOTES ↗ RANK

DOMPurify - a DOM-only, super-fast, uber-tolerant XSS sanitizer for HTML, MathML and SVG. DOMPurify works with a secure default, but offers a lot of configurability and hooks. Demo:

> DOMPurify 3.4.15 focuses on security hardening, adding stronger defenses against DOM clobbering in XML content and fixing several smaller sanitization edge cases.

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

01 **Hardening for XML clobbering and sanitization edge cases**

IMPROVED 30

DOMPurify 3.4.15 adds stronger DOM-clobbering hardening specifically for XML content parsing paths, plus several smaller hardening improvements and edge-case fixes to the sanitization engine.

— Names XML clobbering fix but no specific mechanism, flags, or cases listed. 3.4.15

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

1 RELEASE · 2026-09-05 NOTES ↗ RANK

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 200+ 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's v4.0.1 release adds a URL custom field type, lets teams reopen submitted requirement assessment assignments, and rolls out a batch of smaller UI and workflow refinements alongside findings/ TPRM work.

└──▷ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 **URL type for custom fields** NEW 55

Adds a `URL` type for custom fields, enabling hyperlink values to be stored and displayed on any entity.

– Names the field type and scope but no example of use v4.0.1

02 **Reopen requirement assessment assignments** IMPROVED 50

Adds the ability to reopen a requirement assessment's assignment after it has been submitted, allowing corrections without needing to restart the review process.

– Describes the workflow change but no exact UI steps v4.0.1

03 **Clickable badges on compliance assessments** IMPROVED 45

Makes badges clickable on compliance assessments, enabling faster navigation directly from a badge to related content.

– Names the UI area but not the exact navigation path v4.0.1

04 **Category prefix in control autocomplete** IMPROVED 45

Shows category prefix in control autocompletes, making it clearer which category a control belongs to when searching for controls.

– Small, clearly named UI change but limited detail v4.0.1

05 **Findings and TPRM evolutions** IMPROVED 25

Introduces findings and TPRM (Third-Party Risk Management) evolutions as a major feature epic; the release notes do not detail the specific changes involved.

– No concrete mechanism or scope given v4.0.1

06 **Tasks and evidences improvements** IMPROVED 15

Improves tasks and evidences with additional capabilities, though the specific enhancements are not described in the release notes.

— No specifics on what was added or changed `v4.0.1`

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ai-safe2-framework

1 RELEASE · 2026-09-06 NOTES ↗ RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> ai-safe2-framework's CLI update centers on a new `safe2 doctor` command that discovers agent harnesses, assesses environment posture, and emits Decision Cards, backed by new policy enforcement, schema/evidence, and feedback subcommands — alongside a breaking change to the NEXUS and SkillSpector evidence contracts.

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

01 **Multi-harness discovery and Decision Cards via `safe2 doctor`**

NEW 93

Adds `safe2 doctor` for multi-harness discovery, identifying command and configuration indicators for Codex, Claude Code, Antigravity, Hermes, OpenClaw, and Grok alongside host, shell, CI, and WSL metadata. New `--assess`, `--output`, `--card-format`, and `--card-output` options evaluate environment posture and emit Markdown or HTML Decision Cards containing facts, assumptions, conflicts, history, impacts, alternatives, pros and cons, recommendations, ownership gaps, and exit criteria. Assessment now explicitly targets WSL distributions and SSH-accessible Linux hosts or cloud VMs, surfacing incomplete or unreachable targets as coverage gaps.

— Names exact flags, harnesses and card contents; command is directly runnable

2026-05-09_safe2_CLI_Update

02 **Schema and evidence manifest subcommands** NEW 85

Adds `safe2 schema list`, `safe2 schema validate`, and `safe2 evidence manifest` subcommands for versioned schema discovery, evidence validation, and run manifest assembly with digests, source contracts, integrity status, and coverage.

— Names all three subcommands and the manifest's fields 2026-05-09_safe2_CLI_Update

03 **Versioned evidence contracts for NEXUS and SkillSpector** BREAKING 85

Adds dedicated evidence contracts `safe2.nexus-evidence.v1` and `safe2.skillspector-evidence.v1` for NEXUS and SkillSpector adapters, required by `safe2 evidence manifest`. NEXUS and SkillSpector provider bundles generated before this release must be regenerated to remain compatible with these contracts.

— Names both contract identifiers and the required migration action

2026-05-09_safe2_CLI_Update

04 **Feedback capture via `safe2 feedback`** NEW 81

Adds `safe2 feedback record` and `safe2 feedback summary` subcommands to capture sanitized operational feedback including false completion, missing evidence, silent failures, context loss, stuck loops, permission friction, and integration failures.

— Names both subcommands and all feedback categories captured 2026-05-09_safe2_CLI_Update

05 **Environment policy enforcement with exit codes** NEW 80

Adds environment policy enforcement with `ALLOW`, `DENY`, and `HOLD` outcomes producing exit codes `0`, `1`, and `2` respectively, for use in agent and CI pipelines.

— Exact exit codes make this immediately usable in CI 2026-05-09_safe2_CLI_Update

06 **Config inspection via `--inspect-config`** NEW 78

Adds `--inspect-config` opt-in to `safe2 doctor` to collect an allowlisted structural summary of supported local JSON/TOML configurations, including selected approval and sandbox settings, without emitting raw configuration values.

— Names the flag and its scope, including the no-raw-values guarantee

2026-05-09_safe2_CLI_Update

07 **`safe2 example verify` workflow demo** NEW 60

Adds `safe2 example verify` subcommand to reproduce a baseline-to-drift-to-policy workflow with evidence, Decision Cards, feedback, and manifest validation.

— Names the command but describes the workflow only in brief 2026-05-09_safe2_CLI_Update

08 **Baseline and drift comparison for environments** NEW 60

Adds reviewed baseline and drift comparison across assets, harness indicators, configuration hashes, target coverage, and assessment scope against an earlier inventory.

— Explains comparison scope but no command or flag is named 2026-05-09_safe2_CLI_Update

└──▷ BREAKING ON UPGRADE

! NEXUS and SkillSpector provider bundles generated before this release must be regenerated to be compatible with the new `safe2.nexus-evidence.v1` and `safe2.skillspector-evidence.v1` contracts required by `safe2 evidence manifest`.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 6, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe