THE AI TOOLCHAIN

NO. 080

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 1, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 176 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 54 tools

▾ // THE BRIEF HIDE

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. browser-use's BU Ultrafast is a browser-agent model it says runs over 3x faster and cheaper than the default, picked with one API parameter. Sortie now pauses and resumes

vim ? help

work when a provider's usage limit hits instead of failing the run. magpie defers MCP tool-search loading, cutting the token overhead paid on every request.

BUILD

Let an agent drive a real browser without hand-rolling session plumbing

browser-use's BU Ultrafast swaps in a faster, cheaper browser agent via an API parameter, so the speed/cost tradeoff is a config choice rather than a rewrite; LangGraph's Managed Deep Agents now carry per-thread or per-user browser sessions with human handoff, so a login or a CAPTCHA no longer ends the run. playwright-mcp adds per-session browser context isolation, which stops concurrent agent sessions from leaking cookies and state into each other.

browser-use · LangGraph · playwright-mcp

OPERATE

Stop losing an agent run to a rate limit or a bloated tool manifest

Sortie reads provider usage limits and pauses and resumes work automatically, so hitting a quota becomes a wait rather than a failed job you rerun by hand, and token and cost accounting now spans all agent kinds. magpie defers MCP tool-search loading so large tool catalogs stop being re-sent on every request, and Graphify adds prompt caching on the Claude backend with an environment variable to turn it off.

magpie · Graphify

GOVERN

Keep agent egress and secrets inside a policy you can point at

Brig's route broker puts brokered policy routes in front of outbound traffic, so what a sandboxed workload can reach is a declared rule rather than whatever the network allows, and its secret subsystem now refreshes OAuth tokens instead of expiring mid-run. Arcade replaces raw OAuth URLs with short authorization links and pauses a tool call mid-execution to prompt for input, so an agent asks for consent in-line rather than dying on a missing grant.

Brig

BUILD

◆ AI Coding Agents

repoprompt-ce

Feedback ⓘ v1.8.0 NOTES ↗ RANK

Repoprompt-ce is a native macOS context engineering app for AI coding agents, with an MCP CLI (community edition of RepoPrompt).

> RepoPrompt CE v1.8.0 adds GPT-6 Astra Pro model support, an opt-in Auto Workflow for agent sessions, and performance/UX improvements to the Agent pane, Context Builder, and file cataloging for large workspaces.

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

01 **Agent pane file and image attachments** IMPROVED 50

Files can now be dropped anywhere in an Agent pane to attach them, image attachments are persisted across session restoration, and image previews can be opened directly from the Agent pane.

– Groups three related Agent pane attachment improvements with clear UI actions v1.8.0

02 **Opt-in Auto Workflow for agent sessions** NEW 40

A new opt-in Auto Workflow setting lets newly created agent sessions automatically start Overseer-managed runs.

– Describes behavior but no config path given v1.8.0

03 **Context Builder reviews sibling roots** IMPROVED 40

Context Builder now reviews files across sibling roots, not just the active root, expanding the scope of context it can draw from.

– Explains behavior change but lacks configuration specifics v1.8.0

04 **GPT-6 Astra Pro model support** NEW 35

RepoPrompt CE now supports GPT-6 Astra Pro through the OpenAI API.

– Names model and API but no setup detail v1.8.0

05 **Incremental file-catalog updates** IMPROVED 30

File-catalog updates are now incremental, improving responsiveness when working in large workspaces.

– States the mechanism briefly but no metrics or scope v1.8.0

WAS THIS USEFUL? 👍 👎

yetone magpie

Feedback ⓘ v0.1.2606 CODE ↗ RANK

Magpie is a macOS menu-bar app, TUI, and CLI that lists every AI coding agent on the machine and switches each one's model, such as Codex on DeepSeek or Claude Code on Kimi, by rewriting only that key in the agent's own config.

> Magpie shipped eleven releases focused on Claude Code integration depth and router cost/quota visibility: off-peak pricing for Ollama Cloud, auto-updating Claude Code binaries, deferred MCP tool-search loading that cuts large per-request token overhead, Claude bridge session branching and cache preservation, cross-account quota aggregation, and a breaking change to the browser extension's install location with new drag-to-install support.

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

01 Browser extension install path change and drag-to-install

95

BREAKING

The browser extension is now written to `~/Magpie Browser Extension` (visible in Finder/Explorer) instead of a hidden settings folder; existing installs from the old location keep working since that directory is still kept up to date, but new installs use the new path — this is a breaking change to the install location. On macOS, drag-to-install lets users drag the extension folder from the Library page directly onto `chrome://extensions` without using `Load unpacked`; on Windows, the folder is revealed in Explorer with `/select` so it is ready to drag onto `chrome://extensions`.

— Exact path, UI target, and compatibility note given v0.1.2588

02 Session-scoped usage-log reads via VisitSessions NEW

90

A new VisitSessions(since, sessions, fn) function in `internal/usage/ledger.go` scopes usage-log reads to specific sessions, reading only blocks that name the requested `Session` or `NativeSession` values instead of scanning every record, avoiding full-log scans for lead answer lookups.

— Exact function signature and file path given v0.1.2600

03 ENABLE_TOOL_SEARCH defers MCP tool definitions NEW

90

magpie now sets `ENABLE_TOOL_SEARCH=true` in Claude Code's `~/.claude/settings.json` under `env.ENABLE_TOOL_SEARCH` while Claude Code runs through magpie, deferring MCP tool definitions so they load on demand instead of being sent whole on every request — cutting what was 124K tokens of tool definitions per request for magpie's 245-tool MCP router. It also injects

`ENABLE_TOOL_SEARCH=true` into T3 Code session environments via the `claudeSearchEnv` environment field, covering Claude Code versions before 2.1.304 that kept tool search off behind a non-Anthropic `ANTHROPIC_BASE_URL`.

— Exact config path, token figure, and version threshold `v0.1.2579`

04 **Claude Code interactive-mode environment and caching controls**

IMPROVED 85

Claude Code interactive mode now receives safeguards via `CLAUDE_CODE_EXTRA_BODY` in the run environment. `DISABLE_PROMPT_CACHING` was added to the set of environment variables magpie propagates to Claude Code runs (alongside `ANTHROPIC_BETAS`, `DISABLE_TELEMETRY`, and others in `claudeEnvVars`), enabling prompt-caching control in interactive mode. In interactive mode, client cache marks are now handled separately: words are pasted without marks and Claude Code keeps its own prompt caching, rather than suppressing caching with `DISABLE_PROMPT_CACHING=1` as in non-interactive runs.

— Names exact env vars and caching behavior change `v0.1.2603`

05 **Claude bridge session branching and cache preservation** NEW

85

The Claude bridge now tracks an `allowed` map per run turn so a `tool_choice none` request still offers the client's tools to Claude Code for prompt-cache reuse while preventing any tool calls. It also records `behindLine` and `lastLine` / `backLine` session-line positions so interactive-mode runs can branch off a session at a precise line via a cut-copy fork, replacing the previous `-p` resume that read none of the cached conversation. Side questions are now branched off the current session before being asked, rather than telling the whole conversation in one message, preserving the interactive run's cache across asides.

— Names internal fields and the tool_choice none flag `v0.1.2581`

06 **Combined quota view across accounts and routing groups** NEW

80

The Usage card and tray panel now show an `All {n} accounts` aggregate quota line per provider, summing counted windows by unit, averaging percentage-only windows across accounts on the same plan, and marking failed or stale readings as unknown rather than zero. Routing group rows now display a `groupQuota` summary of remaining allowances across the group's member accounts per provider, refreshed automatically when Usage page readings update unless the group is being edited.

— Names UI surfaces and the groupQuota field `v0.1.2598`

07 **Auto-updating downloaded Claude Code binary** NEW 75

magpie now auto-updates the Claude Code binary it downloaded, checking for the newest release one minute after startup and every six hours after, controlled by the `DISABLE_AUTOUPDATER` environment variable to pin a version.

— Names exact env var and update schedule v0.1.2605

---

08 **Off-peak pricing for Ollama Cloud in router** NEW 70

The cheapest-first router now recognises Ollama Cloud (the `ollama-cloud` preset) as a peak-aware provider, pricing `deepseek-v4.1-flash` and `deepseek-v4-pro` (including dated tags such as `deepseek-v4.1-flash:1017` and `deepseek-v4-pro:0813`) at half-rate outside 12:00–18:00 UTC Monday–Friday, so a relay priced between the two rates is preferred during peak hours and deprioritised off-peak.

— Automatic routing logic; no user action to take v0.1.2606

---

09 **Fallback-model passthrough on retracted refusal** IMPROVED 70

When Claude Code's primary model returns `stop_reason: refusal` and Claude Code retries on its `model_refusal_fallback`, magpie now passes through the fallback model's answer rather than ending the turn with a refusal error.

— Names exact stop_reason and fallback mechanism v0.1.2580

---

10 **Skill market pulls trending skills with pagination** IMPROVED 65

The skill market now pulls skills from both the skills.sh front page and its `/trending` page, so trending skills not already in the front-page list also appear as market offerings, and results are shown 60 at a time with a Show more control to load the next page of cards.

— Names page size and source pages, UI control named v0.1.2585

---

11 **Per-caller browser sessions in MCP router** NEW 65

The MCP router now gives each caller its own Magpie Browser session instead of sharing one across all callers. When Magpie Browser is connected in two browsers simultaneously, an agent can list available browsers and pick one; subsequent calls for that session go to the chosen browser, with each agent working in its own tab and tab group.

— Describes behavior without a named API or flag v0.1.2578

UPDATED INTEGRATION MCP market popular servers snapshot

NEW INTEGRATION Antigravity desktop app (Google) gateway integration

└──▷ **BREAKING ON UPGRADE**

! The browser extension is now written to `~/Magpie Browser Extension` instead of the previous hidden folder inside the settings directory. Browsers that already loaded the extension from the old location continue to work (the old directory is still kept up to date), but new installs use the new path.

v0.1.2587 CODE ↗ STABLE

magpie's MCP market pre-loads the most-starred servers from a shipped snapshot and raises the per-search result limit to 100.

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/yetone/magpie:0.1.2587
```

git

```
$ git clone --branch v0.1.2587 https://github.com/yetone/magpie.git
# already have the repo? check out this version:
$ git checkout v0.1.2587
```

published under the project's GitHub account · checked 1 Nov Wrong command?

MCP market search now returns up to 100 registry results instead of 30, with registry-ranked results shown first followed by matching popular servers.

v0.1.2584 CODE ↗ STABLE

magpie adds an Antigravity (Google desktop app, v2.15+) integration, routing requests through `~/.antigravity/settings.json`; `AGY_LLM_GATEWAY_URL` takes precedence.

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/yetone/magpie:0.1.2584
```

git

```
$ git clone --branch v0.1.2584 https://github.com/yetone/magpie.git
# already have the repo? check out this version:
$ git checkout v0.1.2584
```

published under the project's GitHub account · checked 1 Nov Wrong command?

> Route Antigravity's desktop app requests through magpie's gateway by configuring `~/.antigravity/settings.json` with `wireProtocol` `openai` or `genai`; `AGY_LLM_GATEWAY_URL` overrides the file.
>
> An `agy` CLI session not connected on its own row follows the Antigravity desktop app's gateway row automatically.

WAS THIS USEFUL? 👍 👎

forgedock

Feedback ⓘ v1.41.0 NOTES ↗ RANK

ForgeDock — autonomous AI dev pipeline for Claude Code. GitHub as agent memory: issue in, PR out.

> ForgeDock's biggest addition this window is a `forgedock prod-watch` command for autonomous production monitoring, alongside new lane preflight admission checks, extended conformance calibration, warm sandbox pools for faster lane starts, and a terminal-fitted `forgedock guide` help screen.

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

01 **Production monitoring via prod-watch command** NEW 91

Run production monitoring with `forgedock prod-watch`, a new CLI command with a configurable poll loop, LLM-authored issue writer, dry-run mode, kill switch, owner digest, and P0 paging. Issue volume is controlled with governor settings `max_issues_per_hour`, `max_open_prod_issues`, and a cost cap per round, with an automatic storm summary when the cap is hit. A signal source is automatically dropped to shadow mode when its false-positive rate crosses a threshold, with loop metrics surfaced in `FORGE:TELEMETRY`. Incoming signals are triaged with fingerprinting, deduplication, classification, and validation before any issue is filed.

— Runnable command plus named config keys and telemetry field v1.39.0

02 **Calibrated conformance gating** IMPROVED 73

Adds calibrated conformance gating that measures false-block rates per check kind and enforces only kinds with near-zero false blocks, now extended with mutation guard metrics (survived-mutation share and false-survived counts) and reporting of timed-out validate commands (exit code 124) as `conformance.unimplemented` in lean precheck.

— Mechanism and named status/exit code but no direct command v1.41.0

THINNER COVERAGE BELOW

03 **Lane preflight admission codes** NEW 55

ForgeDock adds lane preflight that refuses admission with typed codes for pool auth failures, pool image version mismatches, PATH shim loops, and missing `gh` field support before a lane starts work.

— Names four failure codes but no reader-facing command v1.41.0

04 **Warm sandbox pools for lane cold-start** NEW 55

Adds warm sandboxes and a same-filesystem dependency cache to cut lane cold-start time, with a per-repo warm sandbox pool managed on the pool agent.

— Explains mechanism but no user-facing control surface named v1.41.0

05 **Terminal-fitted guide command** NEW 55

Run `forgedock guide` to open a new help screen sized to the current terminal.

Get a terminal-fitted overview of forgedock commands before starting a batch run.



— Simple runnable command, thin description v1.40.0

06 **Per-lane phase ledger** NEW 45

Adds a per-lane phase ledger that records where lane wall time goes, plus acceptance-probe coverage metrics.

— Describes what it tracks but not how to view it v1.41.0

WAS THIS USEFUL? 👍 👎

keploy

Feedback ⓘ v3.6.194 NOTES ↗ RANK

Open-source platform for creating safe, isolated production sandboxes for API, integration, and E2E testing.

> Keploy hardened its proxy with a TLS startup gate and added mock value verification to catch drift during test replay.

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

01 **Client TLS gate in proxy** NEW 40

Keploy's proxy now includes a client TLS gate that holds the TLS handshake until the application process is ready, preventing premature connection failures during startup.

— Explains mechanism but no config or flag named v3.6.194

02 **Mock value verification on replay** NEW 40

Keploy's mock replay now validates that the application passes through the exact values a test chose during recording, catching mock value drift on every replay run.

— Describes behavior but no named surface or command v3.6.192

WAS THIS USEFUL? 👍 👎

claude-mem

Feedback ⓘ v15.8.0 NOTES ↗ RANK

claude-mem gives coding agents persistent memory by capturing session activity, compressing it with a model, and injecting relevant context into later sessions.

> claude-mem's sole release this window adds a cloud-hosted viewer for the My Work view and syncs work state across accounts.

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

01 **Cloud viewer and cross-account sync for My Work** NEW

40

claude-mem adds a cloud viewer host for the My Work view and now syncs `work_state` across every account.

— Names the view and work_state key but no mechanism or setup steps v15.8.0

WAS THIS USEFUL? 👍 👎

Alibaba Qwen Code

Feedback ⓘ desktop-v0.26.5 NOTES ↗

RANK

An open-source AI coding agent that lives in your terminal.

> Qwen Code's two 0.26.5 releases add new `--in-place` co-driving modes for `autofix-pr` and `pr-watch`, Kubernetes and Helm deployment options for Managed Agent, QQ chat moderation tools, a batch workspace-snapshot API, and a round of web shell polish including background-agent status, workflow switching, and bounded SQL history queries.

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

01 autofix-pr reliability: stale-base merging, retry budget, human handoff

IMPROVED 88

`/autofix-pr` now merges a stale base before fixing red checks and reports a blocked round with the merge boundary in its status output; it also adds a retry budget so any check still failing after three fix pushes is handed off to a human and surfaced in `/autofix-pr status`, posting a pull-request comment stating what was attempted, what blocks it, and what human action is needed.

— Names exact command, status surface, retry count, and comment content v0.26.5 desktop-v0.26.5

02 `--in-place` co-driving mode for autofix-pr and pr-watch NEW

75

Added `--in-place` mode to both the `autofix-pr` command and the `pr-watch` skill, letting the agent co-drive changes directly in the current working tree or local checkout, push the branch, and post a reply/comment on the pull request.

— Names exact flag and two commands it applies to v0.26.5 desktop-v0.26.5

03 Kubernetes and Helm deployment for Managed Agent NEW

70

Added experimental Kubernetes deployment support for managed-agent server deployments, plus an experimental one-command Helm deployment for Managed Agent that bundles Hosted Harness, MySQL, local execution, and an authenticated gateway.

— Names components but not the exact deploy command v0.26.5 desktop-v0.26.5

04 QQ moderation tools via channel adapters NEW

70

Enabled QQ moderation tools — recall, mute, kick, and join approval — exposed to agent sessions through channel adapters, controlled by the `moderationTools` setting.

— Names config key and the specific moderation actions v0.26.5

05 **Bounded SQL history queries for managed backend** IMPROVED

60

Added bounded SQL history queries to the managed backend to limit query size, covering Managed local H2 and Hosted MySQL, enabling paged recent messages, journal transactions, and exact operation or command lookups.

— Names backends and query types but no runnable example v0.26.5 desktop-v0.26.5

06 **Workflow switch for saved session preferences** NEW

60

Added a Workflow switch in Web Shell, split panes, and VS Code chat that updates saved session Workflow mode preferences for the next turn without interrupting in-progress responses.

— Names three UI surfaces but no config key v0.26.5 desktop-v0.26.5

07 **Web Shell UI polish: paging, shortcuts, locale, context previews**

IMPROVED

60

Web Shell gained live message paging, keyboard shortcuts for Session Overview and Split View, Russian locale for the Goal card and Goals dialog, skill argument hints that stay visible while typing, and clickable context memory files in the Context Usage panel that open read-only previews such as `AGENTS.md`.

— Names several UI surfaces, navigable but not runnable desktop-v0.26.5 v0.26.5

THINNER COVERAGE BELOW

08 **Batch workspace-snapshot API** NEW

55

Added a batch request and SDK method to fetch live-session snapshots for up to 20 selected workspaces in a single call.

— Names the SDK method's limit but not its exact signature v0.26.5

09 **One-click daemon update coverage extended** IMPROVED

50

Extended Web Shell one-click daemon updates to pnpm, yarn, bun, and sudo-prefixed npm installs, and now shows manual update steps when the automatic update is unavailable.

— Names package managers but no exact command v0.26.5 desktop-v0.26.5

10 **Background agent count in status-bar task pill** NEW

40

Web Shell and Qwen Code Desktop now show a count of active background agents in the status-bar task pill while agents are running.

— Thin description of a status indicator `v0.26.5` `desktop-v0.26.5`

11 **Generated UI on by default** IMPROVED

30

Generated UI is now turned on by default so cards stay interactive without manual activation.

— Single-line description of a default change `desktop-v0.26.5`

WAS THIS USEFUL? 👍 👎

Graphify Labs Graphify

Feedback ⓘ v0.9.100 NOTES ↗

RANK

Graphify converts code, docs, SQL schemas, configs, and PDFs into a queryable knowledge graph with explainable relationships.

> Graphify's only update this window adds prompt caching for the Claude backend, with an environment variable to disable it when needed.

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

01 **Claude prompt caching with opt-out flag** NEW

80

Graphify now caches the Claude backend's system prompt using `cache_control`. Gateways that reject this can disable it by setting the environment variable `GRAPHIFY_CLAUDE_PROMPT_CACHE=0`.

Prevent Claude gateway rejections caused by `cache_control` in the system prompt by disabling prompt caching at the environment level.

```
$ GRAPHIFY_CLAUDE_PROMPT_CACHE=0 graphify query "What calls into APIRouter?"
```

— Names exact mechanism and env var with runnable example `v0.9.100`

WAS THIS USEFUL? 👍 👎

coffee-cli

Feedback ⓘ v3.8.3 NOTES ↗ RANK

coffee-cli is a clean AI terminal workspace with quick input, folder switching, copy & paste, and image pasting.

> Coffee CLI's biggest addition this window is DeepSeek Harness as a full desktop conversation agent with multi-agent orchestration, alongside a new OpenCode desktop conversation tab, Claude Code next-prompt predictions, and auto-generated conversation titles shown across tabs and history.

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

01 **Next-prompt predictions in Claude Code** NEW 85

Claude Code desktop conversations now show a predicted next prompt in gray after each reply; pressing Tab fills it in for further editing or sending. It can be disabled by setting `"promptSuggestionEnabled": false` in Claude Code's `settings.json`.

Turn off next-prompt predictions in Claude Code desktop conversations to avoid the extra model request per reply.

```json
{
  "promptSuggestionEnabled": false
}
```

— Names the exact config key, value, and file to disable it. v3.8.3

02 **OpenCode and DeepSeek Harness desktop agent support** NEW 80

Adds desktop agent support for `OpenCode` and `DeepSeek Harness`: `OpenCode` gets a desktop conversation tab with model switching, Build / Plan mode, effort selection, permission approval cards, question answering, launch command selection, usage statistics, computer control, and browser control, and no longer creates an empty OpenCode Session when a tab is opened and closed without sending a message; `DeepSeek Harness` gets a dedicated tab in Coffee CLI with conversation history, model switching, effort-level switching, and balance display, approval and question prompt cards, tab status indicators showing idle, working, and waiting states in green, orange, and blue, works as a multi-agent orchestration worker controllable by Coffee Agent alongside Claude Code and Codex sessions, and gains support in main memory, skills, computer control, browser control, and usage statistics app settings.

— Lists many concrete capabilities but no commands or config. Enumerates every DeepSeek Harness capability and settings area.

v3.8.2 v3.8.3 v3.8.1

03 **Idle-context prompt for stale sessions** NEW 65

Adds an idle-context prompt that asks whether to start a new Claude Code conversation when a session has been idle for over 70 minutes and its context is large, avoiding a full context reload; Codex asks after 12 hours idle.

— Specifies exact idle thresholds and the behavior avoided. v3.8.2

04 **Clickable file paths in conversations and terminal** NEW 65

Adds clickable file paths in conversations and the terminal — hover to underline, click to reveal in File Explorer, right-click to copy the path.

— Names exact hover/click/right-click interactions. v3.8.2

05 **Conversation titles on tabs and in history** IMPROVED 60

Claude Code desktop conversations automatically generate a short title after the first prompt, shown on the tab and in history; tabs for Claude Code, Codex, OpenCode, and DeepSeek Harness now show this title instead of only the folder name, including when resuming from history. History filtering was also improved: system-inserted text is stripped from titles, records with no user input are hidden, and slash-command-only conversations keep the command name as their title.

— Explains mechanism and scope but no exact UI path. v3.8.3

THINNER COVERAGE BELOW

06 **Plan usage and API balance display** IMPROVED 55

Adds plan usage and API balance display to terminal tabs, with usage persisting across tab switches and OpenCode 2 keys now readable; plan usage or balance for the active tool also appears beside the model selector in the composer.

— Names UI locations but no exact config or command. v3.8.2 v3.8.1

07 **Back-to-latest button in terminal tabs** NEW 50

Terminal tabs now have a back-to-latest button, placed at bottom center with a dark fill and accent border whose shape follows the active theme.

— Names placement and styling but no invocation detail. v3.8.3

08 **Slash command descriptions translated** NEW 35

Adds Claude Code slash command descriptions translated into 11 languages.

— States the language count with no further mechanism. v3.8.2

09 **Background tasks shown in Claude Code session view** NEW

25

Claude Code now shows background tasks in the session view.

— Bare statement with no mechanism or scope given. v3.8.1

WAS THIS USEFUL? 👍 👎

Amazon Kiro Crew

Feedback ⓘ v0.11.1 NOTES ↗ RANK

Kiro Crew is an open-source workspace for coordinating AI coding agents.

> Kiro Crew's installer now downloads the built-in KAS engine by default, with a new flag to opt out.

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

01 **`--no-kas` installer flag** NEW 75

The one-line and `install.sh` installers now download the built-in KAS engine by default; passing `--no-kas` skips that download.

— Names exact flag and installer scripts, but no further mechanism detail. v0.11.1

WAS THIS USEFUL? 👍 👎

flow-next

Feedback ⓘ flow-next-v9.5.0 NOTES ↗ RANK

flow-next is a workflow plugin that takes a bug, idea or ticket to a verified pull request with specs, cross-model review, live QA, and receipts

> flow-next 9.5.0 adds an unattended autonomous mode for memory migration, speeds up the deps view dramatically, widens clawpatch compatibility, and replaces `--dry-run` previews with a new `--explain` flag while deprecating preview behaviour on `--auto`.

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

01 Unattended memory-migrate via autonomous mode NEW

83

Running `memory-migrate` with `mode:autonomous` or `FLOW_AUTONOMOUS=1` skips prompts and behaves as `mode:autofix`: it keeps originals and applies mechanical defaults automatically, suitable for CI or other unattended contexts.

Run memory migration in CI or another unattended context without being prompted for input.

```
$ FLOW_AUTONOMOUS=1 /flow-next:memory-migrate
```

— Names exact config key, env var, and resulting behaviour flow-next-v9.5.0

02 New `--explain` flag for route preview NEW

75

`flow --explain <anything>` previews the route flow-next would take for a task without executing any steps, replacing the now-defunct `--dry-run` preview behaviour on `flow --auto`.

Preview the route flow-next would take for a given task without executing any steps.

```
$ /flow-next:flow --explain 'Add passwordless login'
```

— Runnable command shown with exact flag flow-next-v9.5.0

03 `--auto` drops `--dry-run` and `--research` behaviour

65

BREAKING

`flow --auto` no longer accepts `--dry-run` as a preview: it warns, ignores the flag, and runs anyway. It now treats `--research` the same way, warning and continuing to

run instead of honoring it. Use `--explain` to preview a route instead.

— Names both deprecated flags and the replacement `flow-next-v9.5.0`

THINNER COVERAGE BELOW

04 **Faster `/flow-next:deps` view** IMPROVED 57

The `/flow-next:deps` view now calls `flowctl show` only for open specs instead of all specs, cutting calls from 295 to 22 on a representative repository while producing identical output.

— Concrete before/after numbers but no reader action `flow-next-v9.5.0`

05 **clawpatch 0.8 support in `map`** IMPROVED 45

`map` now supports clawpatch 0.8, widening the accepted version range to `>=0.4.0 <0.9.0`.

— Names exact version range but minimal mechanism `flow-next-v9.5.0`

└──▷ **BREAKING ON UPGRADE**

! `flow --auto` no longer accepts `--dry-run` as a preview: it warns, ignores the flag, and runs. Use `--explain` to preview a route instead.

! `flow --auto` now ignores `--research` the same way it ignores `--dry-run`: it warns and continues running.

WAS THIS USEFUL?  

ogcode

Feedback ⓘ v0.56.0 NOTES ↗ RANK

An open-source, local-first AI teammate and agent operating environment for turning intent into completed work—with tools, memory, permissions, and human oversight.

> ogcode's two releases add full-text search over conversation memory (via a new API endpoint and a ⌘K palette), a spreadsheet-querying tool for Excel/CSV/TSV files, and a self-syncing model catalogue with per-model pricing overrides and gateway-rate billing.

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

01 **Full-text search over conversation memory** NEW 91

ogcode adds a `GET /api/session/search` HTTP endpoint that returns ranked results across titles, labels, and turn summaries, plus a ⌘K search palette that queries the same fts5 full-text index (built from `.ogcode/memory`) over title, labels, and body, returning matched turns with their dates and labels; minimum query length is two characters and results page at twenty conversations.

Query conversation memory from a script or external tool when you remember the topic but not the session name.

```
$ curl 'https://<your-ogcode-host>/api/session/search?q=caching'
```

— Endpoint, index source, and query limits all given concretely. v0.56.0

02 **Self-syncing model catalogue** NEW 86

ogcode fetches a shared JSON model catalogue at boot, once a day, and on demand from Settings → General, which shows the last sync time and what changed. The sync endpoint is controlled with the `OGCODE_CATALOG_URL` environment variable, or disabled entirely with `OGCODE_CATALOG_SYNC=false`.

Disable the automatic model catalogue sync in a locked-down environment where outbound fetches are not permitted.

```env
OGCODE_CATALOG_SYNC=false
```

Point ogcode at an internal mirror of the model catalogue — useful when the default endpoint is unreachable from your network.

```text
OGCODE_CATALOG_URL=https://catalog.internal.example.com/ogcode/models.json
```

Trigger an on-demand catalogue sync and confirm it landed without restarting the server.

> 📍 In Settings → General, open the model catalogue card and click `Sync now` to fetch the latest catalogue and review what changed.

— Env vars and UI path given with concrete usage examples. v0.55.0

---

03 **Spreadsheet reading and querying via query_sheet** NEW

85

ogcode adds spreadsheet reading and querying for `.xlsx`, `.xlsm`, `.csv`, and `.tsv` files. The new `query_sheet` tool loads a sheet into an in-memory SQLite table and runs a single read-only `SELECT`, capped at 200 rows and 3 seconds, so large spreadsheets can be aggregated without pulling raw rows into context; plain-delimited files above 256 MB are refused by `query_sheet` with a pointer to the streaming workbook form.

— Tool name, formats, and limits given; no runnable example shown. v0.55.0

---

04 **Per-row model pricing and capability overrides** NEW

76

Any model row — user-added, provider-listed, or catalogue-sourced — now accepts overrides for context window, output ceiling, price per million input and output tokens, cache read and write prices, and image support, stored across restarts via migrations `070` through `072`.

— Migration numbers and override fields named, no UI path given. v0.55.0

THINNER COVERAGE BELOW

---

05 **Expandable tool cards with formatted output** IMPROVED

59

Every tool card in chat — including agent-navigation tools such as `file_map` and `update_intent` — now opens to show its output, and `pdf_index` and `docx_index` render results as a page list (page number then labels) rather than raw JSON.

— Names affected tools and output format change, no command given. v0.55.0

---

06 **OGX usage billed at gateway published rates** IMPROVED

55

OGX token usage on the usage page is now billed at the gateway's published per-model rates; an unpublished rate shows as unpriced rather than free, and a user-set price still takes precedence.

— Describes billing behavior without specific numbers or endpoint. v0.55.0

07 **Three-way agent status on agent board** IMPROVED 45

The agent board now splits agent status three ways — running, idle between turns, and finished — so parked sessions are no longer counted as busy.

— States are named but no UI path or trigger is given. v0.56.0

08 **Flattened housekeeping tool rows** IMPROVED 45

Housekeeping tool rows for `codebase_map`, `file_map`, `compact_context`, and `update_intent` now render as flat one-line records with no chevron and are not clickable or focusable.

— Affected tools named but no further mechanism or path. v0.56.0

WAS THIS USEFUL? 👍 👎

cezar

Feedback ⓘ v0.18.0 NOTES ↗ RANK

Open-source orchestrator and ADE for running Claude Code, Codex, OpenCode, Pi and other AI coding agents in parallel — locally or 24/7 on your own server.

> cezar introduces a guided onboarding flow that gets new and returning users signed in, into a project, and connected to GitHub without touching a terminal.

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

01 **First-run and returning-user onboarding flow** NEW 45

cezar adds an onboarding flow that signs the user in to an agent, picks or creates a project, connects GitHub without a terminal, and starts from a welcome page.

— Describes flow steps but no UI paths or config details v0.18.0

WAS THIS USEFUL? 👍 👎

termide

Feedback ⓘ 0.48.0 NOTES ↗ RANK

All-in-one terminal workspace for desktops and servers: editor with LSP, file manager with SFTP/FTP, terminal, git, database viewer and a coding agent in one zero-config static Rust binary.

> Termide 0.48.0 adds a built-in Perforce plugin with full checkout/diff/revert/submit keybindings and agent tooling, opens up plugin-level version-control and UI extension points, lets AI connection keys live in the encrypted vault, and breaks agent token spend down by contributor and model.

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

01 Built-in Perforce plugin NEW 95

Termide ships a built-in Perforce plugin, enabled via Settings → Plugins or `mode = "on"` under `[plugins.perforce]` in the global `config.toml`, with `mode = "opt-in"` allowing per-project opt-in via `enabled = true` in `.termide/config.toml`. Use `Ctrl+E` on a read-only file to check it out for edit, and `Alt+G` in a non-git project to browse opened files by changelist, with `Enter` to open, `F3` to diff, `Backspace` to revert, `Space` for actions, and `Enter` on a changelist to submit it. The plugin marks depot-changed lines in the gutter the same way git does, and exposes `opened`, `diff`, and `edit` tools to the coding agent.

Enable the Perforce plugin globally so the status bar, gutter marks, and Ctrl+E/Alt+G all work for every project without per-project config.

toml

```toml
[plugins.perforce]
mode = "on"
```

Opt individual projects into Perforce when the global config uses opt-in mode, keeping other projects unaffected.

toml

```toml
[plugins.perforce]
enabled = true
```

— Full config keys, keybindings and agent tools named with runnable examples. 0.48.0

02 Version-control extension points for plugins NEW 75

Plugins can now act as a version-control replacement for non-Perforce systems: implement `files/status` to colour files in the file manager, `editor/base` to mark

gutter changes, `file/will_edit` to intercept `Ctrl+E` checkout/revert, and a command with `"replaces": "open_git_status"` to handle `Alt+G` in non-git projects.

— Exact hook names given but aimed at plugin authors, no runnable example. 0.48.0

03 **AI API keys in encrypted vault** NEW 65

AI connection API keys can now be stored in the encrypted password vault from the AI settings page; keys are stored under the connection name and never written to `config.toml`, with `api_key_env` still taking precedence when set.

— Clear UI path and precedence rule given, no exact config key format. 0.48.0

04 **List, diff, and dialog panels for plugins** NEW 60

Plugins can display a `list` panel (groups of rows with actions, folder-tree layout, and buttons) or a `diff` panel, and open a `window/dialog` made of a message, text field, pick list, checkbox list, and up to six buttons.

— Named UI primitives but no usage example or entry point shown. 0.48.0

05 **Per-contributor token spend breakdown** IMPROVED 60

The agent panel's session info (`F3` or `/usage`) now breaks token spend down by contributor — conversation, subagents, `auto` reviewer, `recall` solver, summaries, `/goal` judge, `/handoff`, and Claude Code internal spend — and by model, cache share, and reasoning output.

— Lists exact contributors and shortcut but no example output shown. 0.48.0

WAS THIS USEFUL?  

octomind

Feedback ⓘ 0.67.2 NOTES ↗ RANK

Open-source AI coding agent and agent runtime: one binary, any model, MCP-native. Runs in terminal, CI, or as a daemon.

> Octomind's WebSocket server gains read-only session observation and now enforces size budgets on replay resumes, which can reject previously-accepted oversized requests.

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

01 **Size budgets enforced on WebSocket replay resumes** BREAKING 50

WebSocket replay resumes now enforce size budgets, rejecting requests that exceed the configured limit. This is a breaking change: existing clients sending oversized replay resumes will have those requests rejected.

– Explains behaviour change but no concrete limit or config key named. 0.67.2

02 **Read-only session observation over WebSocket** NEW 45

The WebSocket server now supports read-only session observation, letting external clients watch a live session without sending input.

– Describes capability but no endpoint or flag name given. 0.67.2

└──▷ **BREAKING ON UPGRADE**

! WebSocket replay budgets are now enforced; existing clients that send oversized replay resumes will have those requests rejected.

WAS THIS USEFUL? 👍 👎

codealta

Feedback ⓘ 3.2.0 NOTES ↗ RANK

Agentic Development Environment

> CodeAlta 3.2.0 introduces adaptive per-turn model selection, per-session context size control and OpenTelemetry tracing for agent runs, alongside LaTeX/chemistry/footnote rendering with a visual equation editor and an MCP server inspection panel.

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

01 Adaptive per-turn model selection NEW

91

An experimental Adaptive mode, configured in Settings › Models › Adaptive or via `altacli run --tier adaptive`, selects a model per turn: a fast tier handles normal turns, an Adaptive recheck on the same model runs for turns that changed code, and escalation to the strongest tier happens only for failed turns.

Run a long coding task on Adaptive so cheap turns handle routine work and only genuinely failing turns escalate to the strongest model, then capture a machine-readable log of every model request for cost analysis.

```
$ altacli run --tier adaptive --format jsonl "Add pagination to the reports endpoint" > run-log.jsonl
```

— Exact flag, settings path and per-turn mechanism given 3.2.0

02 OpenTelemetry tracing for agent runs NEW

90

Send OpenTelemetry traces from CodeAlta Desktop, CodeAlta TUI and `altacli` by setting `OTEL_EXPORTER_OTLP_ENDPOINT`; each run emits a span per session, model request and tool call following generative-AI semantic conventions, with no additional assemblies loaded when no endpoint is configured.

Forward every agent run's model calls and tool calls to a local OTLP collector (e.g. Jaeger) to trace token usage and latency across a session.

```
$ OTEL_EXPORTER_OTLP_ENDPOINT=http://localhost:4317 altacli run "Refactor the auth module"
```

— Named env var, scope of spans, and a runnable command 3.2.0

03 Per-session context size control NEW

90

Set a per-session context size via the `context_size` provider setting or the More options of the Agent chip: 200K or 1M tokens on Claude Code, standard or long-context on Copilot and the OpenAI/Anthropic APIs, or the model maximum on Codex.

Pin a provider to its 1M-token long-context window by default so large-codebase sessions never silently fall back to a smaller context.

```yaml
context_size: 1000000
```

— Named config key with exact per-provider values and example 3.2.0

---

04 **LaTeX, chemistry and footnote rendering with equation editor** NEW 90

Adds LaTeX (`$…$` inline, `$$…$$` displayed, `math` blocks), chemistry (`\ce{…}`), and footnote rendering in messages and Markdown documents, with displayed equations copyable as LaTeX or MathML. An equation editor opened with `Alt+=` supports visual formula entry (`/` for fractions, `^` for superscripts, `sqrt` , `alpha` , or symbol picking) and inserts or replaces equations in the prompt or code editor.

⋮

— Exact syntax, shortcut and editor behaviour, with screenshots 3.2.0

---

05 **JSONL audit log for altacli runs** NEW 70

Write `--format jsonl` with `altacli` to emit a `_codealta/operation` line for every model request, enabling per-request audit and cost tracking.

Run a long coding task on Adaptive so cheap turns handle routine work and only genuinely failing turns escalate to the strongest model, then capture a machine-readable log of every model request for cost analysis.

```
$ altacli run --tier adaptive --format jsonl "Add pagination to the reports endpoint" > run-log.jsonl
```

— Named flag and emitted field but schema undetailed 3.2.0

---

06 **MCP server inspection panel** NEW 65

Click the MCP status above the prompt to inspect every MCP server's state, its defining file, and all its tools with descriptions and agent-callable names; search across all tools by name or description.

— Clear UI path and scope, no command named 3.2.0

THINNER COVERAGE BELOW

07 **Find sessions outside registered projects** NEW 40

Use `altacli alta` to find sessions of folders outside registered projects.

— Bare command name with no further mechanism 3.2.0

WAS THIS USEFUL? 👍 👎

Amazon Kiro

Feedback ⓘ

CLI: Approve for Me, Subcommand Approval Rules, and an ACP Engine Setting (2026-10-29)

BLOG ↗ RANK

Kiro is an agentic development environment that uses spec-driven workflows to plan, build, and maintain software.

> Kiro's CLI release adds configurable ACP engine selection, autonomous approve for me sessions, subcommand-level shell approval rules, tighter V2 Git-config safeguards, and an upgrade command that flags permission regressions.

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

01 Configurable ACP agent engine NEW

85

ACP sessions can be pinned to a specific engine by setting `acp.agentEngine` to `v1`, `v2`, or `v3` in configuration, or overridden per run with the `--agent-engine` flag.

Pin all ACP sessions in a project to the V3 engine without passing a flag every time.

```yaml
acp:
  agentEngine: v3
```

Override the engine for a single ACP run from the CLI, regardless of config.

```
$ kiro-cli acp --agent-engine v2
```

— Names exact config key, values, and flag with runnable examples.

CLI: Approve for Me, Subcommand Approva…

02 'Approve for me' autonomous approval in V3 sessions NEW

70

Turning on Approve for me in interactive V3 sessions via `/permissions` lets Kiro auto-approve routine actions, block risky ones, and ask only when unsure; the choice is stored as a policy preset in the workspace permissions file.

Enable Approve for me in a V3 session so routine actions run without manual confirmation.

> 📍 1. In an interactive V3 session, run `/permissions` and choose Policies. 2. When the page offers approve for me, select it. 3. The policy preset is saved to the

workspace permissions file for this repository only.

— Explains mechanism and storage but navigation is a UI path, not a command.

CLI: Approve for Me, Subcommand Approva…

---

03 **Permission rule migration warnings via agent upgrade** IMPROVED

70

Running `kiro-cli agent upgrade` surfaces warnings when upgraded V3 rules deny, hide, or ask for something V2 previously allowed, with a warning count per file shown in the result table.

— Names exact runnable command and describes its output.

CLI: Approve for Me, Subcommand Approva…

THINNER COVERAGE BELOW

04 **Subcommand-level shell approval rules** NEW 55

Shell commands can now be approved at the subcommand level from the prompt, choosing between an exact command like `git stash pop`, a wildcard `git stash *`, or all of `git *`.

— Clear mechanism with named examples but no runnable command given.

CLI: Approve for Me, Subcommand Approva…

---

05 **V2 approval gate for Git config env vars** IMPROVED 40

V2 sessions now ask for approval before running any shell command that sets Git configuration through an environment variable such as `GIT_CONFIG_GLOBAL`.

— Behavior change described but no steps to invoke or configure it.

CLI: Approve for Me, Subcommand Approva…

WAS THIS USEFUL?  

◆ Model Training & Fine-Tuning

fizgig

Feedback ⓘ v9.0.1 NOTES ↗ RANK

Fizgig — LoRA & Fine-tune Studio for Klein 9B, Krea 2, MiniMax H3, Qwen Image 2.1, Anima & Z-Image Turbo and Ace-Step XL: train, fine-tune, profile, repair and extract LoRAs & LoKR for Image, Video & Music Models

> Fizgig's v9.0.1 release automates model downloads when starting a run, adds Krea 2 support for ComfyUI int8 checkpoints, and expands RunPod pod management with deletion, automatic Helpers-list joining, and file manager credential display.

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

01 **Automatic model download on Start** NEW 65

Pressing Start on a model whose files are not yet downloaded now triggers Download models for me automatically, then starts the run once the files are in place; Klein and Flux 3 files require a Hugging Face token set on the Preferences tab first.

— Names the trigger, prerequisite token, and affected models. v9.0.1

02 **Pods tab management: deletion, auto-join Helpers, and file manager credentials**

NEW 65

The Pods tab now includes a Delete pod action that removes a RunPod pod and its volume disk after a confirmation step. RunPod pods join the Helpers list automatically once their Fizgig instance responds, with no job ticked, and removing one keeps it removed until added back by hand. On a pod, the sidebar also shows the file manager login credentials — user `admin` and the password — above the file manager button.

— Describes mechanism and exact UI elements but only a navigation path, not a command. v9.0.1

03 **Krea 2 support for ComfyUI int8 checkpoints** NEW 60

Krea 2's DiT slots now accept ComfyUI's int8, int8 ConvRot, and other fp8 checkpoints, so a Krea 2 file already on hand for ComfyUI works directly in Fizgig.

— Names exact checkpoint formats supported but no further mechanism given. v9.0.1

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

Arcade

Feedback ⓘ updated · Nov 1, 2026 NOTES ↗ RANK

Arcade provides an actions runtime that lets AI agents securely call external services with authorization and governance controls.

> Arcade's release centers on four headline capabilities — mid-execution tool prompts that pause and resume agent calls, automatic TTL-based refresh of remote MCP catalogs, short authorization links in place of raw OAuth URLs, and self-registering Kubernetes MCP workers — alongside smaller usage-reporting, deployment, and resource-sizing changes, plus two upgrade-breaking database and permission changes.

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

01 Automatic TTL refresh for remote MCP catalogs NEW

80

Remote MCP catalogs now refresh automatically on a configurable interval instead of requiring a manual refresh, set via `engine.remoteCatalogTtl`.

Keep remote MCP tool catalogs fresh automatically so operators never need to trigger a manual refresh.

```yaml
engine:
  remoteCatalogTtl: 300s
```

— Named config key with a runnable example snapshot-20261101

02 Lower default ephemeral storage for deployed servers BREAKING

80

Deployed servers now request `1Gi` of ephemeral storage by default instead of `6Gi`, configurable via `deployments.sizing.ephemeralStorageRequest`.

Reduce ephemeral storage allocation for deployed MCP servers in a resource-constrained cluster.

```yaml
deployments:
  sizing:
    ephemeralStorageRequest: 1Gi
```

— Exact before/after values plus runnable config example `snapshot-20261101`

03 **Short authorization links** NEW 70

Enable short authorization links in place of long OAuth URLs with `engine.enableShortLinks`.

Surface short authorization links to users instead of raw OAuth URLs, and label this deployment in usage analytics.

yaml

```yaml
engine:
  enableShortLinks: true
usageReporting:
  deploymentDisplayName: 'prod-us-east'
features:
  releaseFeed:
    enabled: true
```

— Named config flag with example usage `snapshot-20261101`

04 **Usage reporting and admin notification controls** NEW 70

Usage reporting adds a configurable deployment label via `usageReporting.deploymentDisplayName`, a new Authorization link clicked event, and an admin warning for unsupported or outdated self-hosted installs enabled with `features.releaseFeed.enabled`.

Surface short authorization links to users instead of raw OAuth URLs, and label this deployment in usage analytics.

yaml

```yaml
engine:
  enableShortLinks: true
usageReporting:
  deploymentDisplayName: 'prod-us-east'
features:
  releaseFeed:
    enabled: true
```

— Three named config keys/events, one with an example `snapshot-20261101`

THINNER COVERAGE BELOW

05 **Mid-execution tool prompts (elicitation)** NEW 55

Tools can now prompt the caller for input mid-execution and resume once answered, even across a reconnect, implemented via a new elicitation exchange column.

— Mechanism described but no runnable config or command `snapshot-20261101`

06 **Kubernetes worker self-registration** NEW 45

Kubernetes MCP deployments now self-register as managed workers and are cataloged automatically once running, requiring a new signing-secret column in the engine.

— Describes behaviour but no config example given `snapshot-20261101`

07 **Upgrade-breaking database index and quota changes** BREAKING 45

The one-connection-per-user-and-provider database index is replaced with a compatibility index and then dropped, with rollback not clean; two self-hosted installs now require a quota-permission check before upgrading.

— Impact described but points reader to external upgrade notes `snapshot-20261101`

08 **Arcade Deploy support for custom MCP-protocol servers** NEW 40

Arcade Deploy can now run a custom MCP-protocol server end to end, failing early if a dependency is missing.

— Behaviour stated without concrete command or flag `snapshot-20261101`

09 **Stateless protocol support for outgoing MCP connections** IMPROVED 25

Outgoing MCP connections now support the newer stateless protocol path.

— Single-line addition with no mechanism or example `snapshot-20261101`

NEW INTEGRATION MCP Apps interactive composers for Gmail, Slack, and Outlook Mail

└──▷ BREAKING ON UPGRADE

! The one-connection-per-user-and-provider database index is replaced with a compatibility index; the old index is then dropped. Rollback is not clean — review Upgrade notes before upgrading.

! Two self-hosted installs require a quota-permission check before upgrading — review Upgrade notes.

> announced · Oct 29, 2026 BLOG ↗
>
> Arcade adds MCP Apps support, rendering interactive in-chat UIs with policy-gated controls and native file upload/download via the MCP protocol extension.

> Arcade now supports MCP Apps (the Model Context Protocol interactive-UI extension), letting tools render editable forms, composers, and dashboards inside the chat instead of switching to a browser tab.
>
> File uploads and downloads now work through MCP Apps: an in-chat interface opens the native file picker, accepts drag-and-drop files such as screenshots, and downloads large files straight to disk, bypassing the current MCP protocol gap for file parameters.
>
> Arcade's authorization and governance runtime enforces access policy at the point of interface construction and again at click time, so UI elements such as Send buttons are suppressed when the calling user lacks the corresponding tool permission.
>
> Agent routines that run unattended are not blocked waiting for human input — the interactive interface is available when a user wants to step in, but routines continue without it.

WAS THIS USEFUL? 👍 👎

claw-orchestrator

Feedback ⓘ v7.16.2 NOTES ↗ RANK

Run Claude Code, Codex, Antigravity, Grok Build and OpenCode behind one runtime, and a run only counts as done when checks it didn't write pass. Cross-engine fan-out, councils, session handoff, spend caps, MCP server, ACP agent.

> claw-orchestrator v7.16.2 adds pinned session persistence for agent nodes, fixes Claude goal token accounting to exclude cache reads, and ensures fresh goal threads receive failure context.

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

01 **Pinned session persistence for agent nodes** NEW 70

A new `sessionName` property lets an agent node keep its conversation across multiple visits and nodes instead of restarting the thread each time it's reached.

Keep an agent node's conversation alive across multiple visits so it retains prior context instead of starting a new thread each time.

yaml

```yaml
sessionName: my-agent-session
```

— Names config key and shows exact usage example v7.16.2

02 **Claude goal token accounting corrected** IMPROVED 60

`tokensUsed` and `tokenBudget` for Claude goals now count input, output, and cache writes only; cache reads are excluded from the totals.

— Names exact fields and the before/after accounting change v7.16.2

THINNER COVERAGE BELOW

03 **Failure context on fresh goal threads** IMPROVED 35

Fresh goal threads now receive failure context on their first visit, covering both a revisit whose thread was never set and a goal node taking over after a red verdict.

— Describes behavior but no named surface or command v7.16.2

WAS THIS USEFUL? 👍 👎

HeyGen HyperFrames

Feedback ⓘ v0.8.226 NOTES ↗ RANK

HyperFrames is an open-source HTML-to-video renderer that runs in AI-agent workflows.

> HyperFrames shipped cloud rendering of final cuts for paid plans alongside several Producer and Studio refinements, including passthrough for full-canvas clips, independent control of preview loaders, and clock reporting in playback health checks.

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

01 **Cloud rendering of final cuts** NEW 50

HyperFrames now renders final cuts in the cloud for users on a paid plan, preserving the composition's original frame rate and full quality.

— Names paid-plan gate and quality/frame-rate preservation but no API or command.

v0.8.224

02 **Full-canvas video layer passthrough in Producer** NEW 50

The Producer gains opt-in passthrough for a lone full-canvas video layer, sending the source stream straight through without re-encoding.

— Explains mechanism (no re-encode) but not how to opt in. v0.8.226

03 **Independent toggle for Studio preview loaders** IMPROVED 45

A host app that draws its own loader can now turn off both of Studio's preview loaders independently.

— Clear behavior change, no exact config/flag name given. v0.8.226

04 **Clock reporting in playback health** IMPROVED 23

HyperFrames now reports which clock a film followed in the playback health reading.

— Vague description with no mechanism or usage path. v0.8.226

WAS THIS USEFUL? 👍 👎

sortie

Feedback ⓘ v1.31.0 NOTES ↗ RANK

sortie turns tracker tickets into autonomous agent sessions.

> Sortie's v1.31.0 release adds provider usage-limit awareness that pauses and resumes work automatically, broader token and cost accounting across all agent kinds, and SSH worker agents now get access to Sortie's tool server.

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

01 Token and cost accounting across all agent kinds IMPROVED

75

`sortie stats` now shows the source and completeness of each token figure for cancelled, killed, and resumed turns across all agent kinds. Sortie now counts token usage and cost for every agent kind on local and SSH worker runs, marking figures as lower bounds or unknown when a run cannot be fully measured instead of showing zero. Codex cost accounting now includes prompt-cache writes, and OpenCode runs now include the token spend of their subagents.

— Names command and specific provider behaviors; no exact flags v1.31.0

02 Provider usage-limit awareness and issue parking NEW

70

Sortie now waits for a provider usage limit to reset instead of retrying, continues sending work to other agents and models in the meantime, and surfaces each agent's plan limits, usage, and reset time on the Overview page. The Overview page also lists issues being held back and the reason each one waits, and releases a parked issue back into the agent's earlier session when the agent supports session resumption.

— Explains mechanism and UI surface but no exact config key v1.31.0

THINNER COVERAGE BELOW

03 Token limits now apply to more agent kinds BREAKING

57

Token and cost limits, including `agent.max_tokens`, now apply to agent kinds they did not count before — runs that previously ran without hitting a limit may now stop sooner with the same settings.

— Names config key and behavior change, aiding migration v1.31.0

04 SSH worker agent tool access NEW

40

Agents running on SSH worker hosts now have access to Sortie's tools and HTTP tool server connections defined in the workflow.

— States the capability but no config or command detail v1.31.0

└──▷ **BREAKING ON UPGRADE**

! Token and cost limits (including `agent.max_tokens`) now apply to agent kinds they did not count before — runs that previously ran without hitting a limit may now stop sooner with the same settings.

WAS THIS USEFUL? 👍 👎

cotal

Feedback ⓘ v0.132.5 NOTES ↗ RANK

The open standard for agent coordination

> Cotal v0.132.5 improves error visibility for embedders and operators, exposing a typed lease-conflict error from the Manager API and surfacing the auth daemon's typed startup refusal in CLI output.

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

01 Typed ManagerInstanceLeaseHeld error from Manager.start()

NEW 85

Manager.start() now throws the exported `ManagerInstanceLeaseHeld` typed error instead of a generic failure, carrying the held lease row and a `sameRoot` flag so embedders can inspect lease conflicts programmatically rather than parsing message strings.

Catch a lease conflict in an embedder without fragile message parsing, using the new `ManagerInstanceLeaseHeld` type and its `sameRoot` field to decide recovery strategy.

javascript

```
import { Manager, ManagerInstanceLeaseHeld } from 'cotal-ai';

try {
  await manager.start();
} catch (err) {
  if (err instanceof ManagerInstanceLeaseHeld) {
    console.error('Lease held by another process. Same root:',
err.sameRoot);
    console.error('Held lease row:', err.lease);
  } else {
    throw err;
  }
}
```

— Runnable code example with exact type and fields named v0.132.5

THINNER COVERAGE BELOW

02 Typed auth daemon refusal shown in cotal up output IMPROVED

50

`cotal up` now surfaces the auth daemon's typed startup refusal in its output, naming the error code and reason (e.g. `[unavailable; ai.cotal.auth.plane-claim-refused: unknown]`) instead of a bare exit.

— Describes output format but no command walkthrough v0.132.5

WAS THIS USEFUL? 👍 👎

Brig

Feedback ⓘ v0.5.0 NOTES ↗ RANK

Brig runs coding agents in a microVM sandbox, made up of the brig CLI and the brigd session daemon.

> Brig's v0.5.0 introduces a route broker with brokered policy routes for structured egress control, makes overlay-root microVM storage the default, adds OAuth token refresh to the secret subsystem, and removes a large set of legacy CLI subcommands, flags, and config keys.

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

01 **Route broker and brokered policy routes** NEW 90

Adds the `v1alpha1` policy envelope, brokered routes, and a route mode registry to the policy system for structured egress control. Adds `brig-broker`, a route broker injector that brokers API routes to the guest microVM and enforces that tokens cannot be used outside their carrier, and routes `brig.internal` to microVMs so guests can reach brokered endpoints by hostname. `brig info` now shows broker status alongside the rest of the execution envelope.

— Names envelope version, injector, hostname and CLI surface, but no config example

v0.5.0

02 **Removed legacy CLI subcommands, flags and config keys** 85

BREAKING

The subcommands `brig profiles`, `brig agents`, `brig profile`, `brig template`, `brig import`, `brig export`, `brig policies`, `brig create`, `brig env`, `brig reset`, `agent list`, `agent save`, `agent load`, `policy list`, `secret list` and `secret rm` are removed. The run-line flags `-t`, `-m`, `-n`, `--name`, `-w`, `--workspace`, and `-q` (after the verb) are removed. The profile keys `forward:`, `statePaths:`, `shell:` and `gui:` are removed; a profile that still contains any of them will not parse. The environment variable `BRIG_TEMPLATE_DIR` is removed; Brig refuses to start while it is set and `BRIG_PROFILE_DIR` is not set.

— Enumerates every removed surface verbatim with migration conditions v0.5.0

03 **Overlay-root microVM storage as default** NEW 80

Brig now boots overlay roots by default on microVMs. Filesystem layering can be controlled with the `BRIG_ROOTFS_MODE` environment variable to run an overlay root on two disks, and `brig doctor` now reports whether an overlay root is available on the host.

— Names env var and diagnostic command but not full mode list v0.5.0

04 **OAuth token refresh in secret subsystem** NEW 60

Adds broker and renewer seams to the secret subsystem, with automatic OAuth token refresh for Claude and a journal of each refresh written to the store.

— Describes mechanism but no config key or command to inspect it v0.5.0

THINNER COVERAGE BELOW

05 **Herd profile for running agents** NEW 30

Brig adds a `herd` profile for running agents.

— Bare name with no behaviour or usage detail v0.5.0

└──▷ **BREAKING ON UPGRADE**

! The subcommands `brig profiles`, `brig agents`, `brig profile`, `brig template`, `brig import`, `brig export`, `brig policies`, `brig create`, `brig env` and `brig reset` are removed.

! The subcommands `agent list`, `agent save`, `agent load`, `policy list`, and `secret list` and `secret rm` are removed.

! The run-line flags `-t`, `-m`, `-n`, `--name`, `-w`, `--workspace`, and `-q` (after the verb) are removed.

! The profile keys `forward:`, `statePaths:`, `shell:` and `gui:` are removed; a profile that still contains any of them will not parse.

! The environment variable `BRIG_TEMPLATE_DIR` is removed; Brig refuses to start while it is set and `BRIG_PROFILE_DIR` is not set.

WAS THIS USEFUL? 👍 👎

platypus

Feedback ⓘ v3.30.0 NOTES ↗ RANK

Self-hosted AI Agents for your whole team — on your infrastructure, your models, around the clock.

> platypus's v3.30.0 release focuses on context management, adding per-model reasoning controls, automatic chat summarisation, and a manual `/compact` command to free up context space.

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

01 **/compact command to condense chats** NEW 55

A new `/compact` command lets users condense a chat on demand to free up context space without losing continuity.

— Names exact command but no detail on mechanism v3.30.0

02 **Per-model Reasoning controls** NEW 55

platypus adds Reasoning controls driven by a per-model Reasoning map, which is sent to native providers through the SDK `reasoning` option.

— Names config surface and SDK option but scope unclear v3.30.0

03 **Automatic context management and summarisation** NEW 50

platypus now automatically summarises a chat's older transcript into a durable summary when the context grows large, and attempts to recover from a context overflow, notifying the user when it cannot. It also clears older write tools' input automatically to free up context space, shows a summarising status in the chat transcript while this is in progress, and displays a UI hint when no context window is declared for a model.

— Describes multiple behaviours but no config surfaces or limits named v3.30.0

WAS THIS USEFUL? 👍 👎

persona

Feedback ⓘ @runtypelabs/persona@4.34.0 NOTES ↗ RANK

Create agentic front-end experiences for the web in VanillaJS. Lightweight, extensible, and WebMCP-native.

> Persona 4.34.0 adds collapsible handling for long pasted or transcribed user messages and a new Approval needed indicator on the jump-to-latest control.

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

01 **Collapsible long user messages** NEW 88

Long user messages (pasted text, voice transcripts) can now be collapsed behind Show more / Show less buttons by enabling `features.userMessageCollapse`. The collapse height is set via `maxHeight` (default 240px), labels are customized with `copy.showMoreLabel` and `copy.showLessLabel`, and the fade/toggle appearance is styled through `theme.components.userMessageCollapse` tokens. It is off by default but enabled under `future.v5Defaults`.

Prevent long pasted or voice-transcribed messages from filling the viewport above an assistant reply in anchor-top mode.

javascript

```
createPersona({
  features: {
    userMessageCollapse: {
      maxHeight: 240
    }
  },
  copy: {
    showMoreLabel: 'Show more',
    showLessLabel: 'Show less'
  },
  theme: {
    components: {
      userMessageCollapse: {
        // your custom tokens here
      }
    }
  }
});
```

— Names exact config keys, default value, and runnable example.

@runtypelabs/persona@4.34.0

02 **Approval needed indicator in jump-to-latest control** NEW

The jump-to-latest control now shows an `Approval needed` pill when a pending approval is scrolled below the fold. The label is configured via `features.scrollToBottom.approvalLabel` and styled with `theme.components.scrollToBottom.approval` tokens (`background`, `foreground`, `border`, `fontWeight`); custom `renderIcon` implementations can detect this state via `state.activity === 'approval-below'`.

Surface an `Approval needed` pill in the jump-to-latest control and style it to match your brand when an approval gate is off-screen.

javascript

```
createPersona({
  features: {
    scrollToBottom: {
      approvalLabel: 'Approval needed'
    }
  },
  theme: {
    components: {
      scrollToBottom: {
        approval: {
          background: '#1e40af',
          foreground: '#ffffff',
          border: '#1e3a8a',
          fontWeight: '600'
        }
      }
    }
  }
});
```

— Specifies config keys, style tokens, and state flag with example.

`@runtypelabs/persona@4.34.0`

WAS THIS USEFUL? 👍 👎

browser-use

Feedback ⓘ announced · Oct 31, 2026 BLOG ↗ RANK

The browser-use Python library automates web browsers for AI agents.

> browser-use launched BU Ultrafast, a new browser agent model that completes tasks over 3x faster and cheaper than the default agent, selectable via a simple API parameter alongside a standard-cost variant.

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

01 BU Ultrafast and BU Fast agent models NEW 90

A new `bu-ultrafast` browser agent model can be selected on all Cloud plans by setting `"model": "bu-ultrafast"` in the `POST /api/v4/runs` request body, with no other changes required; it runs at 2x priority token billing. The same agent is available at standard (1x) billing via `"model": "bu-fast"`. BU Ultrafast completes tasks in a median of 22 seconds at $0.0064 per task, versus 79 seconds and $0.0098 for the default GPT-6 Luna agent — over 3x faster and cheaper.

Run a web task with the new Ultrafast agent to cut median task time from 79 s to 22 s at lower cost.

```
$ curl https://api.browser-use.com/api/v4/runs \
  -H "X-Browser-Use-API-Key: $BROWSER_USE_API_KEY" \
  -H "Content-Type: application/json" \
  -d '{"task": "Find the cheapest flight from SFO to JFK next
Friday", "model": "bu-ultrafast"}'
```

Use bu-fast when the task needs more reasoning per step and you want standard (1x) token billing instead of the 2x priority rate.

```
$ curl https://api.browser-use.com/api/v4/runs \
  -H "X-Browser-Use-API-Key: $BROWSER_USE_API_KEY" \
  -H "Content-Type: application/json" \
  -d '{"task": "Research and summarise the top 5 competitors of
Acme Corp", "model": "bu-fast"}'
```

— Names endpoint, exact model values, and concrete speed/cost figures.

launch-20261031-518d44b0

WAS THIS USEFUL? 👍 👎

LangChain LangGraph

Feedback ⓘ announced · Oct 29, 2026

BLOG ↗ RANK

Low-level orchestration framework for building, managing, and deploying long-running, stateful AI agents.

> LangGraph's Managed Deep Agents gained a managed browser subsystem, giving agents automated web-UI control with per-thread or per-user session persistence and human handoff.

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

01 **Managed browser for Deep Agents** NEW 91

Agents can be given a managed browser by exporting a `browser` from `browser.py` using `define_browser` from `managed_deepagents`, with `scope` set to `'thread'` (default, isolated browser per conversation) or `'user'` (one browser shared across a signed-in user's threads, persisting cookies and login state) and `idle_ttl_seconds` controlling the idle timeout. Managed Deep Agents automatically adds four tools: `browser_snapshot` reads the page accessibility tree, `browser_run` executes batched actions or Playwright-style JavaScript via Stagehand, `browser_screenshot` captures the rendered page, and `browser_handoff` pauses the agent at a point like a login page so a human can sign in interactively in LangSmith Studio before resuming, with credentials persisted for later threads under `scope='user'`. The browser runs in a separate LangSmith sandbox connected via a private localhost bridge to the agent's sandbox, enabling end-to-end testing of apps on `localhost` ports without a public URL.

Give a Managed Deep Agent a per-thread browser with a 10-minute idle timeout so it can click through web UIs that have no API.

python

```python
# browser.py
from managed_deepagents import define_browser

browser = define_browser(
    scope="thread",
    idle_ttl_seconds=600,
)
```

Share one browser across all of a signed-in user's threads so the agent only asks for a login once and reuses the saved session in later conversations.

```python
# browser.py
from managed_deepagents import define_browser

browser = define_browser(
    scope="user",
    idle_ttl_seconds=1800,
)
```

— Names config function, flags, four tools, and runnable code examples.

launch-20261029-14b0d726

WAS THIS USEFUL? 👍 👎

◆ MCP Servers & Tooling

elementor-mcp

Feedback ⓘ v3.21.1 NOTES ↗ RANK

WordPress MCP server with 234 Free tools and 547 total with Pro for building pages, managing content, and operating your site. Availability depends on installed integrations and enabled tools.

> elementor-mcp's v3.21.1 release adds ten native plugin integrations across free and Pro tiers and introduces a safety framework for write tools with capability checks, confirmation, and conflict-aware undo history.

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

01 **Ten native plugin integrations added** NEW 70

Six free-tier integrations are now available in all tiers: Redirection, Custom Post Type UI, Pods, WP Mail SMTP, LiteSpeed Cache, and Wordfence, each operating through its plugin's own API and verified against specific plugin versions. The Pro tier additionally gains FluentSMTP, Flamingo, Site Kit by Google, and Ultimate Member.

— Names all ten integrations but no setup steps given v3.21.1

THINNER COVERAGE BELOW

02 **Safety framework for write tools** NEW 55

New write tools start disabled and check capabilities, confirmation, plugin constraints, and supported versions before executing. Reversible writes are logged to History with conflict-aware undo, while irreversible actions describe what cannot be recovered.

— Explains mechanism well but no concrete flags or commands v3.21.1

WAS THIS USEFUL? 👍 👎

claudexor

Feedback ⓘ v5.0.1 NOTES ↗ RANK

Multi-harness control plane for Claude Code, Codex, Cursor, and OpenCode: quota-aware rotation across multiple Claude/Codex subscriptions, shared thread context, and cross-model review.

> Claudexor's v5.0.1 release pins PR delivery to a required Betterleaks 1.9.0 scanner, replaces the apply endpoint's soft-failure response with a typed HTTP 409 contract, and removes the size ceiling on Git-based result storage, while also retiring credential-shaped-text masking and the `review --delta-scope` flag.

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

01 Betterleaks 1.9.0 required for PR delivery

BREAKING

90

Claudexor's `pr` delivery now requires the pinned open-source scanner Betterleaks 1.9.0. An npm-installed CLI bundles no scanner and refuses with the typed error `publication_scan_unavailable` until it is configured, either by setting `CLAUDEXOR_BETTERLEAKS_PATH` to a Betterleaks 1.9.0 build or by running `node scripts/provision-betterleaks.mjs` to download and pin it.

Point Claudexor at a local Betterleaks 1.9.0 build so `pr` delivery can scan before publication.

```
$ export CLAUDEXOR_BETTERLEAKS_PATH=/opt/betterleaks-
1.9.0/betterleaks
claudexor doctor
```

Let the built-in provisioner download and pin the required Betterleaks 1.9.0 scanner instead of pointing to a manual build.

```
$ node scripts/provision-betterleaks.mjs
```

— names exact env var, provisioning script, version, and error code v5.0.1

02 Typed HTTP 409 contract on apply endpoint

BREAKING

90

A `commit`, `branch`, or `pr` delivery that fails after writing now answers `POST /v2/runs/:id/apply` with a typed HTTP 409 problem response and

`accept_clean_patch`, replacing the previous HTTP 200 `applied:false` response.

Handle the new HTTP 409 typed failure when a delivery fails after writing, instead of checking for HTTP 200 `applied:false`.

```
$ curl -s -X POST https://localhost:PORT/v2/runs/:id/apply \
  -H 'Content-Type: application/json' \
  -d '{"accept_clean_patch": true}'
```

— exact endpoint, method, and before/after status codes given v5.0.1

THINNER COVERAGE BELOW

03 **Credential masking removed from user work** BREAKING 45

Credential-shaped text in user work is no longer refused or masked, and the typed error `inline_secret_rejected` is removed.

— names removed error code but gives no replacement behavior v5.0.1

04 **`review --delta-scope` flag retired** DEPRECATED 40

`claudexor review --delta-scope` is retired and no longer accepted.

— names exact flag but no rationale or replacement given v5.0.1

05 **Unlimited-size Git result storage** IMPROVED 35

Large run results are now stored as Git objects with no size ceiling, removing the previous limit on result size.

— mechanism named but no concrete limit or number given v5.0.1

└──▷ **BREAKING ON UPGRADE**

! `pr` delivery now requires Betterleaks 1.9.0; an npm-installed CLI refuses with `publication_scan_unavailable` until `CLAUDEXOR_BETTERLEAKS_PATH` points at the pinned build or `node scripts/provision-betterleaks.mjs` provisions it.

! Credential-shaped text in user work is no longer refused or masked, and the typed error `inline_secret_rejected` is removed.

! `claudexor review --delta-scope` is retired and no longer accepted.

! `POST /v2/runs/:id/apply` now answers HTTP 409 with a typed code and `accept_clean_patch` when a `commit`, `branch`, or `pr` delivery fails after writing; the previous response was HTTP 200 `applied:false`.

WAS THIS USEFUL? 👍 👎

open-codex-computer-use

Feedback ⓘ v1.1.3 NOTES ↗

RANK

Open Computer Use – Open-Source Alternative to Codex Computer Use

> Open Computer Use's v1.1.3 release adds macOS text selection and cross-platform accessibility-tree/screenshot observation options, tightens automation accuracy and performance across macOS, Windows, and Linux, and makes virtual displays an explicit opt-in rather than a default.

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

01 **Virtual displays now require explicit opt-in** BREAKING

55

Virtual displays no longer activate automatically — they now require explicit opt-in, so any setup that relied on virtual displays being on by default will stop working after upgrading.

– Clear before/after behavior change but no exact opt-in flag named v1.1.3

02 **Cross-platform automation accuracy and performance improvements**

IMPROVED

52

macOS automation gains more accurate coordinate clicks, focused-sheet snapshots, live-focus text replacement, and faster accessibility reads; Windows gets improved scrolling and scaled-display input handling; Linux sees reduced observation overhead.

– Lists specific improvements per platform but no metrics or config v1.1.3

03 **Optional accessibility-tree or screenshot collection** NEW

40

Adds optional accessibility-tree or screenshot collection across platforms, letting callers choose their observation method per session.

– Describes behavior but not the config/flag to select it v1.1.3

04 **macOS text selection for agents** NEW

35

Adds macOS text selection capability, giving agents a new interaction surface for selecting and manipulating text on screen.

– Names capability but no mechanism or API detail v1.1.3

└──▷ **BREAKING ON UPGRADE**

! Virtual displays no longer activate automatically — they now require explicit opt-in, so any setup that relied on virtual displays being on by default will stop working after upgrading.

WAS THIS USEFUL? 👍 👎

kordoc

Feedback ⓘ v4.24.34 NOTES ↗ RANK

모두 파싱해버리겠다 — HWP·HWPX·PDF·Office 문서를 Markdown으로. 양식 자동 채우기와 신구대조를 갖춘 CLI·MCP 서버 | Convert Korean documents (HWP, HWPX, PDF, Office) to Markdown — CLI and MCP server with form filling and diff

> kordoc's `to-hwpx` command now supports batch conversion across multiple files and entire folder trees, not just single files.

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

01 **Batch folder conversion in `to-hwpx`** IMPROVED 93

The `to-hwpx` command now accepts multiple files and folders for batch HWP-to-HWPX conversion: folders are scanned recursively for `.hwp` and `.hwpml` files, each converted to `.hwpx` alongside the original. Existing `.hwpx` files are skipped, conversion continues past broken files, and a summary of converted/skipped/failed counts is printed with exit code 1 if any failures occur; the `-o` flag remains available for single-file use only.

Batch-convert an entire directory of legacy HWP files to HWPX, skipping already-converted files and getting a pass/fail summary for CI.

```
$ npx kordoc to-hwpx ./문서폴더
```

— Exact flag, behaviour, exit code and skip logic all specified. v4.24.34

WAS THIS USEFUL? 👍 👎

davinci-resolve-mcp

Feedback ⓘ v4.18.0 NOTES ↗ RANK

davinci-resolve-mcp is an MCP server integration for DaVinci Resolve Studio.

> davinci-resolve-mcp adds a pair of live caption commands that let callers read and edit captions in an open DaVinci Resolve project without restarting the application, including safe rollback on write failures.

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

01 **Live caption writing with destructive-action safeguards** NEW

90

The new `timeline_ai.write_captions_live` command writes, adds, or deletes captions in an open project, accepting the same `replace` / `add` / `delete` shape as the offline `project_db` `write_captions` command. It exports a native DRT, applies the edits, and imports a new timeline revision, leaving the original timeline unchanged if any step fails. It is registered as a destructive action with a native dry run, classified as medium risk and recoverable.

— Full mechanism, risk classification, and rollback behavior given v4.18.0

02 **Live caption reading for open timelines** NEW 70

The new `timeline_ai.list_captions_live` command reads captions directly from an open timeline, returning caption text, per-word timing, and stable `id` values (the start frame as a string) that are safe to pass to the writer command. It is registered as a read-only action.

— Names exact command and output fields, but no limits or edge cases v4.18.0

WAS THIS USEFUL? 👍 👎

playwright-mcp

Feedback ⓘ v0.0.84 NOTES ↗ RANK

Playwright MCP server

> playwright-mcp v0.0.84 adds per-session browser context isolation, runtime URL parameter redaction, bounded find results, cross-navigation network history, cross-language locator support, and finer-grained cookie deletion.

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

01 **Redaction of runtime URL parameters** NEW 90

Use `--redact-url-params=<param>` (config key `redactUrlParams`, env `PLAYWRIGHT_MCP_REDACT_URL_PARAMS`) to replace named query and fragment parameter values with `<redacted>` in tool responses, useful for stripping tokens and auth codes before they reach MCP logs.

Strip OAuth tokens and auth codes from tool responses so they never appear in MCP logs.

```
$ npx @playwright/mcp@0.0.84 --redact-url-
params=access_token,code,id_token
```

— Flag, config key, env var and runnable example all given v0.0.84

02 **Isolated browser contexts per session** NEW 88

Pass `isolatedContext: "<name>"` to `browser_tabs` with `action: "new"` to isolate sessions, e.g. by account: tabs sharing a name share a context, different names stay isolated, and omitting the parameter or using `"default"` uses the default context.

— Exact tool, action and field with full mechanism given v0.0.84

03 **Cross-navigation network request history** NEW 85

Setting `all: true` on `browser_network_requests` retains up to 200 requests from earlier page navigations with stable indexes, so they can be inspected later via `browser_network_request`.

— Named fields and a concrete limit of 200 requests v0.0.84

04 **Skip workspace-root negotiation** NEW 85

Use `--ignore-roots` (config key `ignoreRoots: true`, env `PLAYWRIGHT_MCP_IGNORE_ROOTS=true`) to skip unsupported workspace-root negotiation, falling back to the server's current working directory as the workspace root.

— Flag, config key and env var fully named v0.0.84

05 **Filtered cookie deletion** IMPROVED 78

`browser_cookie_delete` (opt-in via `--caps=storage`) gains optional `domain` and `path` filter fields, letting cookies be deleted by name, domain, and path.

— Tool, capability flag and new fields all named v0.0.84

06 **Bounded browser_find results** NEW 70

The new `maxResults` integer field on `browser_find` returns only the first N matching items along with their context, while still reporting the total match count.

— Named field explained but no runnable command shown v0.0.84

07 **Cross-language locators in element tools** IMPROVED 65

Element tools such as `browser_click` now accept JavaScript, Python, Java, or C# locator strings in the `target` field, alongside existing snapshot refs and CSS selectors.

— Names the field and supported languages, no example v0.0.84

WAS THIS USEFUL? 👍 👎

3dsmax-mcp

Feedback ⓘ v1.7.7 NOTES ↗ RANK

3ds Max MCP server

> 3dsmax-mcp v1.7.7 adds animation capture over multiple frames with viewpoint/shading control, async job polling for long-running calls, and new diagnostics for busy or hung 3ds Max instances, plus a session-pinning safeguard and material-replacement tweaks.

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

01 Async job polling for render and capture calls NEW

90

`render_scene` and `capture_multi_view` return async job handles after `wait_seconds` (maximum 50) elapses while 3ds Max keeps working; passing `job=` to the same tool later collects the result without re-sending the work.

Kick off a long render without blocking; poll for the result once 3ds Max finishes.

```
$ # First call — returns a job handle after wait_seconds
result = render_scene(wait_seconds=50)
job_id = result["job"]

# Later call — collect the finished render
final = render_scene(job=job_id)
```

— Named parameters plus a runnable example v1.7.7

02 Diagnostics for busy or hung 3ds Max instances NEW

90

New `MAX_BUSY` and `MAX_NOT_RESPONDING` status codes surface instance state, `get_bridge_status` responds even when Max is hung and reports what is holding it, and `capture_hang_diagnostics` reads where Max is stuck without contacting it.

Check why 3ds Max is unresponsive without contacting Max directly, then read where it is stuck.

```
$ status = get_bridge_status()
# If MAX_NOT_RESPONDING, read the hang details
diag = capture_hang_diagnostics()
```

— Named status codes and tools with a runnable example v1.7.7

03 **Multi-frame animation capture with viewpoint and shading control**

NEW

80

`capture_multi_view` now captures animation previews as a single tiled image using `frames` or `frame_count` parameters, holding the camera framed over the subject so motion reads across tiles — covering rigs, simulations, and tyFlow. Capture viewpoint can be set with `three_quarter`, `current`, or any `[x, y, z]` direction, and `shading` can be set to smooth, flat, wireframe, or clay with `edges` for topology visibility, applied per capture only.

— Names all parameters but lacks a runnable example v1.7.7

THINNER COVERAGE BELOW

04 **Session pinning via expected_pid** NEW

55

Passing `expected_pid` to any tool refuses a call that would reach a different 3ds Max instance than the one the session opened.

— Names the flag but no usage example v1.7.7

05 **Material replacement source control and Revit bitmap reading**

IMPROVED

55

`replace_material` behavior is updated with `source_from` to control where source material comes from while leaving sub-materials untouched; material graphs now also read Autodesk Bitmap paths from Revit materials.

— Names the parameter and behavior but no example v1.7.7

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ AI / LLM Gateways

Mozilla AI Otari

Feedback ⓘ v0.35.0 NOTES ↗ RANK

Open-source, OpenAI-compatible LLM gateway you run yourself. One endpoint for 40+ providers, with virtual keys, budgets, and usage tracking.

> Mozilla AI Otari ships a dedicated subagent router API and dashboard, adding token-price-aware routing recommendations and free (zero-price) router configs, while moving router configuration exclusively into the dashboard.

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

01 **Subagent router API and dashboard** NEW 76

Subagent routers can now be managed through the new `/routing/recommenders` API endpoint, which stores and uses token prices when making routing recommendations and supports zero-price (free) configurations. A new subagent router page was added to the web dashboard to configure and manage these routers, including setting router name and token prices (which may be 0).

Navigate to the subagent router page in the web dashboard to configure a router without the `config.yml` model option.

📍 1. Open the Otari dashboard. 2. Go to Routing › Subagent Router. 3. Click Add Router, fill in the router name and token prices (both may be 0 for free models). 4. Save — the router is now active for recommendations.

– Names endpoint and dashboard steps but lacks pricing mechanics detail. v0.35.0

THINNER COVERAGE BELOW

02 **Dashboard-only subagent router configuration** BREAKING 35

Subagent routers can now only be set via the dashboard; any previous method of setting the subagent router outside the dashboard no longer works.

– States the change but gives no migration path or old-method details. v0.35.0

└──▷ BREAKING ON UPGRADE

! Subagent routers can now only be set via the dashboard; any previous method of setting the subagent router outside the dashboard will no longer work.

WAS THIS USEFUL? 👍 👎

◆ Inference Providers & GPU Clouds

Eigen Labs Darkbloom

Feedback ⓘ v0.10.6 NOTES ↗

CODE ↗ RANK

Darkbloom is a for the innovators pushing the frontier.

> Darkbloom's v0.10.6 release enables tensor-parallel serving across a pair of Apple Silicon Macs over Thunderbolt 5 without a coordinator, alongside a cleaner CLI download progress display.

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

01 Two-Mac tensor-parallel cluster serving over Thunderbolt 5

NEW 83

Darkbloom can now serve models like GLM-5.3-Flash and Qwen3.8-27B in tensor parallel across a pair of Apple Silicon Macs connected via Thunderbolt 5. Running `darkbloom cluster local` serves the pair without a coordinator, keeping the pair on the leading Mac's local API.

Start tensor-parallel cluster serving on two Macs joined by Thunderbolt 5 without needing a remote coordinator.

```
$ darkbloom cluster local
```

— Names command, hardware link and models, but no scaling limits detailed v0.10.6

THINNER COVERAGE BELOW

02 Single-line model download progress in CLI IMPROVED

25

The CLI now shows model download progress on one updating line instead of printing a stream of separate lines.

— Describes the change but no command or config to act on v0.10.6

WAS THIS USEFUL? 👍 👎

◆ Agentic Infrastructure

sandboxd

Feedback ⓘ v0.3.26 NOTES ↗ RANK

Open-source, self-hosted AI app builder — an agent builds real apps in isolated sandboxes on your own server, each live at a preview URL. Self-host in one command.

> sandboxd adds OpenClaw as a second agent on the Agents page, giving users a Hermes alternative for building and modifying apps under the same safety limits.

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

01 **OpenClaw agent on Agents page** NEW 75

sandboxd adds OpenClaw alongside Hermes on the Agents page: connect it with a model, chat with it, and ask it to build or modify apps. It cannot delete apps, is capped at up to 10 new apps per day, runs under 3 GB of memory from its official image, and carries the same protections as Hermes.

Add OpenClaw on the Agents page to let it build or change apps on your behalf, without risking deletion of existing ones.

📍 In the sandboxd UI, go to Agents › OpenClaw, supply your model connection, then chat with OpenClaw to request app creation or changes. The page shows OpenClaw's download size and memory use (up to 3 GB) before you confirm.

— Names limits and mechanism; UI path given but no API/flag. v0.3.26

WAS THIS USEFUL? 👍 👎

◆ Inference Serving & Optimization

vLLM

Feedback ⓘ announced · Oct 29, 2026 BLOG ↗

The vLLM inference engine runs large language models with high-throughput, memory-efficient GPU execution.

vLLM adds Mistral Large 4 serving with DSpark speculative decoding, decode/prefill context parallelism, cross-layer KV sharing, and multimodal image input support.

NEW SUPPORTED PLATFORM NEW INTEGRATION Serve Mistral Large 4 via `vllm serve` with FP8 or NVFP4 checkpoints, using `--tool-call-parser mistral`, `--enable-auto-tool-choice`, and `--reasoning-parser mistral`.

Enable DSpark speculative decoding with `--speculative-config '{"method": "dspark", "num_speculative_tokens": 3}'` for approximately 2x output throughput on TP8.

Shard the latent KV cache for long-context workloads using `--decode-context-parallel-size 8`, raising KV capacity up to 7.9x on TP8.

Allow multiple images per prompt with `--limit-mm-per-prompt '{"image": N}'`.

Cross-layer KV sharing added via `kv_sharing_target_layer_name` on MLAAttention and `AttentionLayerBase`, available model-agnostically.

WAS THIS USEFUL? 👍 👎

DATA

◆ Vector DBs & RAG

oxideterm

Feedback ⓘ v2.3.0 NOTES ↗ RANK

AI-native workspace for local shells and remote machines.Zero Webview, zero OpenSSL, zero telemetry, and no app subscription.Written in Rust+GPUI

> OxideTerm v2.3.0 adds BLE terminal connections and expands Kitty/Sixel graphics support, while cutting terminal history memory by up to 92% through text/style separation and tighter cache management, alongside SFTP, SSH session, and Markdown preview improvements.

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

01 BLE terminal connections NEW

88

OxideTerm adds a BLE terminal that scans Bluetooth adapters and nearby devices, inspects GATT services, and connects to the Nordic UART Service or custom services with compatible notification/write characteristics. Connections can be saved, cloud-synced, and included in encrypted `.oxide` imports and exports, and support text and hexadecimal send/receive with separate input/output line-ending settings and configurable local echo, MTU, and write mode.

Connect to a Bluetooth Low Energy device running Nordic UART Service to send and receive serial data without a USB cable.

> 📍 In OxideTerm, open New Connection, select BLE, scan for nearby devices, choose the target device from the GATT service list, set line endings and write mode as needed, then click Connect.

— Named surfaces and step-by-step connect flow given v2.3.0

02 Terminal memory and image budget reductions IMPROVED

80

Terminal history memory drops from 196.47 MB to 15.50 MB (92.1% reduction) for 20,000 full ASCII history rows at 400 columns by storing text and styles separately and sharing styles across adjacent characters. Search reduces peak extra heap for CJK and combining-character regex hits from 119.79 MiB to 74.13 MiB by sharing text snapshots and using smaller position records. Hidden terminal tabs release glyph, paint, highlighting, and link-layout caches, cutting additional heap allocations from 60.02 MiB to 35.25 MiB in an eight-terminal, seven-hidden-tab benchmark. The per-terminal image retention budget at default quality is now 256 MiB, with terminal images, decode/upload buffers, and background images sharing a 512 MiB CPU image budget.

.

— Precise before/after benchmarks and mechanism, no user action needed v2.3.0

03 **SFTP sidebar views and sync** NEW 75

The SFTP sidebar now lets you switch between list and directory-tree views, expand subdirectories on demand, and remember your view preference. SFTP also supports syncing the selected remote directory to the corresponding terminal, dragging files into an SSH terminal to upload to its current directory, and path bookmarks with rename, delete, and per-connection management.

Upload files from your local machine to a remote SSH session's current working directory by dragging them directly into the terminal.

> 📍 In OxideTerm, open an active SSH terminal session, then drag one or more files from your local file manager into the terminal window to upload them via the corresponding SFTP connection.

— Concrete UI actions with a drag-and-drop upload example v2.3.0

04 **Kitty and Sixel graphics improvements** IMPROVED 73

Kitty graphics now support Unicode placeholders, virtual placements, and relative positioning so images follow placeholder text or a parent image across main/alternate screen switches, partial scrolling, reflow, and history trimming. Kitty animations now support frame composition, partial updates, transparency, loop counts, pause/resume, loading waits, and frame deletion, with playback reusing complete decoded frames. Sixel images now follow text through scrolling, reflow, and cleanup, clearing only the affected image area on text overwrite, and keep native size inside tmux panes.

— Detailed mechanism but no user-facing command or UI path v2.3.0

THINNER COVERAGE BELOW

05 **Markdown preview in IDE** NEW 45

Markdown files in the IDE now open in preview by default, with mouse selection, cross-paragraph copy, and one-click switch back to source while preserving focus, unsaved edits, and undo history.

— Clear behavior but no exact UI path or shortcut named v2.3.0

06 **SSH session manager grouping** IMPROVED 35

Active SSH sessions in the session manager are now grouped by host and user, show saved connection names, and distinguish port and jump-host paths.

— Describes change but no steps or config given v2.3.0

WAS THIS USEFUL? 👍 👎

1c-ai-connector

Feedback ⓘ v1.9.0 NOTES ↗ RANK

ИИкона — AI-коннектор для 1С:Предприятие 8.3: LLM-провайдеры, агентская петля (function calling), RAG, MCP-сервер, мониторинг ошибок с ИИ-анализом, генератор диаграмм

> 1c-ai-connector v1.9.0 adds multi-conversation tabs with independent histories and extends the `get_object_structure` method with query-table field inspection, though the latter requires synchronized upgrades across databases.

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

01 query_table parameter for get_object_structure

BREAKING

85

`get_object_structure` gains an optional `query_table` parameter that returns real fields, types, and parameter names — including virtual register tables — for a selected query-language table, with parameters listed in call order. This is a breaking change: a remote extension that ignores `query_table` will no longer return ordinary object attributes as table fields, so both the central and target databases must be updated to v1.9.0 together.

— Names exact parameter, method, and migration requirement v1.9.0

THINNER COVERAGE BELOW

02 Multi-conversation tabs in a single form NEW

58

Users can run multiple saved conversations with different experts and models within one form, each with its own history, draft, materials, and settings. Background requests continue to save responses to the originating conversation even after switching tabs.

— Describes mechanism and scope but no UI path given v1.9.0

└──▷ **BREAKING ON UPGRADE**

! A remote extension that ignores `query_table` will no longer return ordinary object attributes as table fields — both the central and target databases must be updated to v1.9.0 together.

WAS THIS USEFUL? 👍 👎

liliang-cn CortexDB

Feedback ⓘ v2.142.0 NOTES ↗ RANK

AI memory and a knowledge graph in one SQLite file. Pure Go: vectors, RAG, agent memory, RDF/SPARQL, Cypher, 80+ MCP tools.

> CortexDB added per-operation control over its vector index backend and introduced a chunk-filtering option for graph node queries.

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

01 **skip_chunks option for graph node matching** NEW 75

Passing `skip_chunks` to `find_nodes` excludes chunks from graph node matches, so queries return only top-level nodes when chunk noise is unwanted.

— Exact parameter and function named, directly usable v2.142.0

02 **Selectable vector index backend per operation** NEW 70

Operators can choose the vector index backend per operation over gRPC — `hnsw` or `binary` on SQLite — allowing speed-versus-recall tuning without rebuilding the database.

— Names backends and transport but not exact param/flag syntax v2.142.0

WAS THIS USEFUL? 👍 👎

EVALUATE

◆ AI Observability & Evals

LangChain LangSmith

Feedback ⓘ updated · Nov 1, 2026

NOTES ↗

LangSmith provides tracing, evaluation, and deployment tools for LLM applications.

LangSmith adds MCP tools for querying runs, managing prompts, sandbox file access, and deployment log inspection, plus OAuth token unification and LSU-based pricing display.

New `query_runs` MCP tool accepts the same parameters as client.runs.query(), returns token counts and USD costs by default, and supports cursor pagination; `fetch_runs` is deprecated in its favor.

The MCP's `get_prompt_by_name` now returns the prompt template itself, and new MCP tools let agents push prompt versions, browse commits, manage commit tags, update or delete prompts, and read and commit Context Hub agent and skill files.

The LangSmith MCP can now read, write, and search sandbox files; stream long-running sandbox commands; share sandbox ports behind LangSmith login; and report sandbox usage, organization usage, audit logs, usage limits, trace retention settings, and organization members.

Sandbox proxy rules now accept `match_headers` entries so several rules on one host can each swap a different placeholder credential and serve multiple accounts.

Agents on LangSmith's native MCP server can find deployments and revisions and read build or runtime logs filtered by revision, time, search text, and severity.

+5 more

NEW INTEGRATION OpenAI model cost resolution expanded to all pricing page models

NEW INTEGRATION Vertex AI Model-as-a-Service open model cost resolution

NEW SUPPORTED PLATFORM New LLM models available in Playground and model configurations

WAS THIS USEFUL? 👍 👎

UK AI Security Institute Inspect AI

Feedback ⓘ

0.3.287 · CODE ↗ · RANK

Inspect is the UK AI Security Institute's framework for large language model evaluations, covering datasets, solvers, scorers, and agentic eval tasks.

> Inspect AI 0.3.287 adds runtime-level rewriting of react() conversations and exposes new budget and token metrics to the sentinel host, alongside a minor registry reservation for upcoming range goal types.

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

01 Runtime conversation rewriting at turn boundaries NEW

60

Runtimes can now hold and rewrite a react() conversation at turn boundaries, and compaction handlers gain a reset() method.

– Names react() and reset() but no full usage example 0.3.287

THINNER COVERAGE BELOW

02 Sentinel host exposes input budget and monitor token counts

IMPROVED

40

The sentinel host now exposes the input budget and the token count of a monitor model to evaluation code.

– States new data exposed but no API detail 0.3.287

03 Reserved range_goal registry type NEW

35

Reserves the `range_goal` registry type for `inspect_ranges` goal types.

– Bare naming of an internal reservation, no usage shown 0.3.287

WAS THIS USEFUL? 👍 👎

PromptLayer

Feedback ⓘ October 31, 2026 NOTES ↗ RANK

PromptLayer is a platform for managing, evaluating, and observing prompts and LLM applications.

> PromptLayer's October 31 update adds structured prompt template version export, new statistical scorecard evaluation metrics, and several Smart Table registry improvements including duplication, renaming, and snippet-version traceability.

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

01 **Prompt template version export** NEW 75

Export a prompt template version as a structured file from the version selector, including `model settings`, `labels`, `input variables`, `tool configuration`, and referenced `snippet versions`.

Take a tested prompt template version offline or into another system, preserving all model settings, variables, and referenced snippets.

📍 In PromptLayer, open the prompt template, open the version selector, and choose Export to download the structured file including model settings, labels, input variables, tool configuration, and referenced snippet versions.

— Lists exact exported fields and exact navigation steps. October 31, 2026

02 **Smart Table duplication from registry** NEW 70

Duplicate an existing Smart Table from its actions menu, choosing whether the copy includes computed values and setting a custom title for the new table.

Reuse an existing Smart Table structure as a starting point for a new experiment without rebuilding columns from scratch.

📍 In PromptLayer, open the Smart Tables registry, click the actions menu on the table you want to copy, select Duplicate, choose whether to include computed values, and enter a custom title for the copy.

— Clear options and step-by-step usage example provided. October 31, 2026

03 **Scorecard statistical evaluation metrics** NEW 65

Scorecard column evaluators now support `Percentile`, `MAE`, and `RMSE` aggregations, plus a per-class classification report for accuracy-based comparison of outputs against expected values.

— Names exact metrics but gives no usage steps. October 31, 2026

THINNER COVERAGE BELOW

04 **Snippet version tracking in Smart Table cells** NEW 45

Smart Table prompt cells now record which snippet versions were used to produce their output, for easier traceability of results back to specific snippet versions.

— Describes behavior but gives no usage path. October 31, 2026

05 **Smart Table rename from registry** IMPROVED 30

Rename Smart Tables directly from the registry menu.

— Single-line description with no further mechanism given. October 31, 2026

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

cplt

Feedback ⓘ 2026.10.31-181434-5e06c82 NOTES ↗ RANK

cplt is a kernel-level sandbox for AI coding agents, running Copilot CLI, Claude Code, OpenCode, Gemini CLI, and others with git and gh guards.

> cplt's `doctor` command now reports on the pnpm sandbox copy, warning when it's missing.

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

01 **pnpm sandbox check in doctor command** IMPROVED 50

The `cplt doctor` command now names the pnpm sandbox copy in its output and warns when no pnpm sandbox copy is present.

— Names exact command and behavior but no mechanism detail 2026.10.31-181434-5e06c82

WAS THIS USEFUL? 👍 👎

containarium

Feedback ⓘ v0.123.0 NOTES ↗ RANK

Open-source agent runtime — SSH-native isolation, eBPF egress policy, Kubernetes + LXC backends, GPU passthrough, MCP-native CLI

> Containarium's latest release adds a no-engine shell command mode, static peer targeting for its cluster, per-member model overrides for crew/graph runs, a signed shared toolchains archive, two platform housekeeping commands, and removes deprecated bridge-DNS flags.

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

01 **Static peer move targets via --peers** NEW 80

Pass static peers as move targets by supplying `id=host:port` to the `--peers` flag, avoiding a dynamic discovery handshake.

Point a static cluster peer as a move target without a dynamic discovery handshake.

```
$ containarium --server http://your-host:8080 --peers
node1=10.0.0.5:9000
```

— Flag, format and a runnable command example are all given. v0.123.0

02 **Shared signed toolchains archive for boxes** NEW 70

Share pinned read-only build tools across boxes via the new managed toolchains archive (`toolchains-linux-<arch>.tar.gz`), distributed with an SPDX SBOM and cosign signature.

— Archive naming, SBOM and signing named but no walkthrough. v0.123.0

03 **Per-member model overrides in crew/graph runs** NEW 65

Override the model per member of a crew or graph run with the new `model_overrides` argument on the `run_crew` tool, propagated to all graph nodes.

— Argument and tool named but no usage example. v0.123.0

THINNER COVERAGE BELOW

04 **StartBoxRun no-engine command mode** NEW 45

Run a shell command as the run user with no engine via the new `StartBoxRun` command mode.

— Names the mode but gives no example or further mechanism. v0.123.0

05 **Sweep stale run checkouts** NEW 45

Sweep stale per-run checkouts of platform runs with the new `sweep-checkouts` code command.

— Command named but behaviour and scope only briefly stated. v0.123.0

06 **Append to existing head branch via SubmitTrackerChangeFiles** NEW 40

Append to an existing head branch and reuse its open change via `SubmitTrackerChangeFiles`.

— Named surface but no mechanism detail or example. v0.123.0

07 **Removed deprecated bridge-DNS flags** BREAKING 40

Deprecated bridge-DNS flags have been removed (refs #2735, #3273); any configuration or scripts that pass those flags will break on upgrade.

— Breaking change noted but exact flag names not given. v0.123.0

└──▷ ALSO FROM THESE RELEASES

Wire the platform MCP server into Claude Code so AI agents can create and manage containers.

json

```
{
  "mcpServers": {
    "containarium": {
      "command": "/usr/local/bin/mcp-server",
      "env": {
        "CONTAINARIUM_SERVER_URL": "http://your-host:8080",
        "CONTAINARIUM_JWT_TOKEN": "<your-token>"
      }
    }
  }
}
```

└──▷ BREAKING ON UPGRADE

! Deprecated bridge-DNS flags have been removed (refs #2735, #3273) — any configuration or scripts that pass those flags will break on upgrade.

WAS THIS USEFUL? 👍 👎

AI MODELS

◆ Open-Weight Models

DeepSeek V3.2 Exp

Feedback ⓘ price cut · Oct 31, 2026

MODEL ↗

DeepSeek V3.2 Exp is a chat model from DeepSeek.

Price cut on OpenRouter (deepseek/deepseek-v3.2): $0.28 → $0.26 in, $0.42 → $0.38 out per million tokens (-10%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

DeepSeek V4 Pro 0813

Feedback ⓘ

price increase · Oct 31, 2026 MODEL ↗

DeepSeek V4 Pro 0813 is a chat model from DeepSeek.

Price increase on OpenRouter (deepseek/deepseek-v4-pro): $0.9553 → $1.60 in, $1.9105 → $3.20 out per million tokens (+67%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

DeepSeek V4.1 Flash

Feedback ⓘ price cut · Oct 31, 2026

MODEL ↗

DeepSeek V4.1 Flash is a multimodal model from DeepSeek.

Price cut on OpenRouter (deepseek/deepseek-v4.1-flash-1017): $0.30 → $0.15 in, $1.20 → $0.60 out per million tokens (-50%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

Google Gemma 4 26B A4B

Feedback ⓘ

price cut · Oct 31, 2026 MODEL ↗

Gemma 4 26B A4B is a multimodal model from Google.

Price cut on OpenRouter (google/gemma-4-26b-a4b-it): $0.09 → $0.0675 in, $0.30 → $0.225 out per million tokens (-25%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

Moonshot AI Kimi K2.7 Code

Feedback ⓘ

price increase · Oct 31, 2026 MODEL ↗

Kimi K2.7 Code is a coding model from Moonshot AI.

Price increase on OpenRouter (moonshotai/kimi-k2.7-code): $0.56 → $0.60 in, $3.10 → $3.50 out per million tokens (+13%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

MiniMax M2.5

Feedback ⓘ price cut · Oct 31, 2026 MODEL ↗

MiniMax M2.5 is a reasoning model from MiniMax.

Price cut on OpenRouter (minimax/minimax-m2.5): $0.30 → $0.27 in, $1.20 → $0.95 out per million tokens (-21%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

◆ Frontier Models

Particle AI Particle-1 Feedback ⓘ

price cut · Oct 31, 2026 MODEL ↗

Particle-1 is a multimodal model from Particle AI.

Price cut on OpenRouter (particle-ai/particle-1): $0.25 → $0.05 in, $1.00 → $0.50 out per million tokens (-80%), first seen 31 Oct 2026.

WAS THIS USEFUL? 👍 👎

Mistral AI Feedback ⓘ announced · Oct 29, 2026 BLOG ↗ RANK

Learn how to deploy and use Mistral AI's Large Language Models with our comprehensive documentation, guides, and tutorials.

> Mistral AI launched Mistral Large 4 (Le Chonk), an open-weight model adding selectable reasoning modes and new defensive cybersecurity capabilities.

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

01 **Defensive cybersecurity task handling** NEW 50

Mistral Large 4 now handles defensive cybersecurity tasks — vulnerability discovery, reproduction, and patching — scoring 51% on the Artificial Analysis Cyber Index while refusing 96.6% of malicious cyber requests.

— Concrete benchmark numbers but no usage instructions launch-20261029-87d5b92d

02 **Three reasoning modes in Mistral Large 4** NEW 45

Mistral Large 4 supports three reasoning modes — Low, Medium, and High (default) — letting developers trade inference cost against task complexity.

— Names the three modes but no API parameter or flag given launch-20261029-87d5b92d

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue