All the cloud native news that's fit to reconcile

The Daily Reconcile.

COUNTDOWN
3 days
to the FOSDEM main track CFP

VOL. I · NO. 317 FRIDAY, NOVEMBER 13, 2026 DAILY EDITION · KUBERNETES & PLATFORM ENGINEERING

INSIDE TODAY

◆ KERNEL SECURITY
An AF_UNIX flaw gives local root, with a public exploit P.2

◆ SUPPLY CHAIN
4,837 hijacked GitHub accounts harvested Actions and cloud secrets P.2

◆ NEW LLMS
Google's Gemini 4 Hydrogen succeeds Flash at $1.50 in, $7.50 out P.2

GITOPS AT SCALE

LinkedIn drops Argo CD for app deploys, and failures fall from 925 to 63

Across more than 250 clusters and 100,000 deployments a week, one pod per deployment replaced a shared controller. Argo CD stays for add-ons.

COMPILED FROM INFOQ, NOVEMBER 12

LinkedIn runs more than 4,000 services on 250+ clusters of up to 5,000 nodes. With one Argo CD instance per cluster, 200+ instances made every upgrade a multi-week project, and a sync took under 3 minutes at 60,000 pods but over 10 minutes past 80,000.

Status was unreliable too: Argo CD has no observedGeneration, a stale cache could report Synced on the wrong revision, and managedFields truncates past 10 field managers, so removed containers could survive. Flux was rejected for the same controller model. Now each deployment runs in its own short-lived pod, uses PUT instead of PATCH and fails on real drift.

From May 2025 to May 2026, deploy-infrastructure failures fell from 925 to 63 while volume rose 3.7 times. Drift errors

Continued on page 2

DEPLOYS AT LINKEDIN, MINUTES

16	DETECTING A FAILED DEPLOY, 2025
2	DETECTING A FAILED DEPLOY, 2026
10	ARGO CD SYNC, 80,000+ PODS
3	ARGO CD SYNC, 60,000 PODS

Median time to detect and conclude a failed deploy fell from 16.1 to 2.45 minutes; sync times are over 10 and under 3 minutes on one cluster. Source: InfoQ, November 12.

AI INFERENCE

Google's TPU Sync cuts a 235B model's KV-cache transfer from seconds to 230 ms

The Apache 2.0 C++ library moves KV-cache blocks and weights between TPU memory, host DRAM, storage and the network, and plugs into vLLM, SGLang and llm-d.

COMPILED FROM INFOQ, NOVEMBER 12

For prefill/decode disaggregation of Qwen3 235B, median KV transfer fell from several seconds to about 230 ms. Average time to first token at low concurrency fell 60% for the 235B model and 90% for a 480B model; prefix caching cut it by up to 90% on 64K-token contexts. All figures are Google's, validated on TPU v7x and v6e.

It ships a KV connector for vLLM-TorchTPU and a backend for SGLang-TorchTPU, and pairs with llm-d 1.0 on GKE for prefix-cache-aware routing.

It is early: the JAX and PyTorch wheels on PyPI are development builds, the PyTorch wheel needs a torch_tpu build that is not public, and the storage tier and RL weight sync are experimental. Pin a known-good commit, as the project advises.

Full story → www.infoq.com/news/2026/11/google-tpu-sync

CONTAINER RUNTIMES

runc, crun and CRIU fix a restore flaw that mounts host paths

COMPILED FROM GITHUB SECURITY ADVISORIES, NOVEMBER 12

A crafted image can make a checkpoint restore bind-mount a host path read-write (three CVEs, 7.9). Fixed in runc 1.3.7, 1.4.4 and 1.5.3, crun 1.31 and CRIU 4.3.1.

Full story → github.com/opencontainers/runc/security/advisories/GHSA-g2c3-j8g3-7gpw github.com/checkpoint-restore/criu/security/advisories/GHSA-7ww7-53q2-wrwx

DEVELOPER PORTALS

Backstage 1.56.5 closes a TechDocs code-execution hole

COMPILED FROM GITHUB ADVISORIES AND RELEASE NOTES, NOVEMBER 12

A user's mkdocs.yml could run code as root in the Zensical build container (7.7); a scaffolder policy bypass rates 8.1. Update the Zensical module to 0.1.1 and scaffolder-backend to 4.2.2.

Full story → github.com/backstage/backstage/security/advisories/GHSA-232c-9pxf-2gp6 github.com/backstage/backstage/releases/tag/v1.56.5

OBSERVABILITY

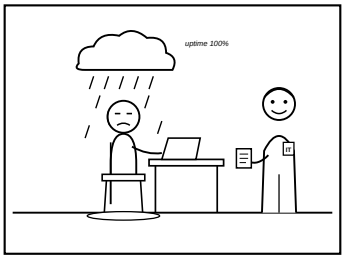
OpenTelemetry deprecates the Collector's batch processor

COMPILED FROM OPENTELEMETRY BLOG, NOVEMBER 11

Exporter batching becomes the default in v0.171.0 (February 2027) and the processor goes in v0.177.0 (May 2027). Drop batch from pipelines and raise queue_size.

Full story → opentelemetry.io/blog/2026/collector-batching-migration

CARTOON OF THE DAY



Drawing: Claude

"Good news: IT says your cloud is highly available."

WHY IT'S FUNNY A small rain cloud has parked itself over one employee's desk, and the IT colleague reports, proudly, that its uptime is 100%. In IT a cloud that never goes down is the dream; over your own chair it is a soggy morning.

DATES TO WATCH

Events and calls for papers

Nov 16 CFP
FOSDEM 2027 main track (Brussels, Jan 30-31): talks due

Nov 16 ONLINE
GitHub pull_request_target brownout, 16:00-18:00 UTC

Nov 17 EVENT
OpenInfra Day North America, Chicago

NEWSROOM

PRESS RELEASE · FROM GITHUB RELEASES AND NOTIFICATIONS

Krateo retires autopilot and tightens identity and CRD checks

November 13, 2026 — On November 12 the installer dropped krateo-autopilot, authn and audit shipped breaking changes and oasgen-provider began enforcing not constraints.

KRATEO PLATFORMOPS

GitHub releases, github.com/krateo-platformops

installer 0.3.596–0.3.610 · Nov 12

- Removes krateo-autopilot; opt-in openstack feature, off by default.

github.com/krateo-platformops/installer/releases

authn, audit 0.43.0–0.45.0, 0.9.0–0.13.0 · Nov 12

- Breaking: no identity in logs; authn 0.45.0 needs frontend 1.8.41.

github.com/krateo-platformops/authn/releases

oasgen-provider 0.61.0–0.63.0 · Nov 12

- Breaking: generated CRDs enforce not constraints (0.63.0).

github.com/krateo-platformops/oasgen-provider/releases

portal, frontend 1.8.110–1.8.113, 1.8.41 · Nov 12

- A table of firing alerts with no incident; per-row actions.

github.com/krateo-platformops/portal/releases

alert-provider 0.2.76–0.2.82 · Nov 12

- Out-of-memory fixes: Incidents load in pages of 100.

github.com/krateo-platformops/alert-provider/releases

Also new: observability 3.1.91–3.1.92, otel-collector 1.0.7–1.0.9, core-provider 2.13.86, provider-runtime 1.17.1–1.18.0, unstructured-runtime 1.13.0–1.15.0, git-provider 0.17.0, nightly-review 0.2.0–0.2.2 (Nov 8–12). Latest elsewhere: snowplow 1.12.70 (Nov 11), helm-render-service 0.5.0 (Sep 30).

KRATEO AGENTIKO

Merged into main this week

Krateo Agentiko notifications could not be read this morning.

CONTINUED FROM PAGE 1 · GITOPS AT SCALE

LinkedIn GitOps

are excluded, and LinkedIn credits orchestrator work too. The engine is internal, not open source; Argo CD stays for add-ons. Check your own setup: count field managers on long-lived objects and see whether sync time tracks total pods.

Full story → www.infoq.com/news/2026/11/kubernetes-argocd-linkedin

NEWS IN BRIEF

KERNEL SECURITY

An AF_UNIX flaw gives local root, with a public exploit

CVE-2026-80521 (7.8) affects kernels 6.10 to 7.1.9 and some 6.1 and 6.6 backports; fixes are in 7.1.10, 6.18.53, 6.12.111, 6.6.158 and 6.1.189. Google's public exploit worked in 19 of 20 boots without user namespaces or BPF. Patch shared nodes and CI runners first.

— Daily CyberSecurity, November 12
securityonline.info/cve-2026-80521-linux-kernel-privilege-escalation

REGISTRIES

Red Hat Quay flaws rated 9.8 and 9.9 open every private repository

CVE-2026-123710 lets an anonymous token plant a signing key and push anywhere; CVE-2026-110697 leaks robot tokens. Security firm Ostorlab lists fixes including 3.18.3; Red Hat lists none yet. Block /keys at the proxy.

— Red Hat; Ostorlab, November 11
blog.ostorlab.co/red-hat-quay-registry-takeover-cve-2026-123710.html

SUPPLY CHAIN

4,837 hijacked GitHub accounts harvested Actions and cloud secrets

Since November 10 the wfx-loot campaign probed 61,764 repositories and took secrets from at least 1,835, per Aikido. Look for a wfx-docs branch and "docs: dependency audit notes" pull requests; rotate what workflows could reach.

— BleepingComputer, November 12
www.bleepingcomputer.com/news/security/hackers-compromise-4-800-github-accounts-to-steal-cloud-credentials

DATABASES

PostgreSQL 18.7 and four other lines fix 52 CVEs, and 14 reaches end of life

Most let a logged-in user run code as the server's OS user; none rates 9.0 or more. Update, then reindex suspect GIN, pg_trgm GiST and bpchar B-tree indexes.

— PostgreSQL; Daily CyberSecurity, November 13
www.postgresql.org/support/security

NEW LLMS

Google's Gemini 4 Hydrogen succeeds Flash at $1.50 in, $7.50 out

It has a 1M-token context, is GA in the Gemini API and keeps weights closed. Artificial Analysis scores it 51, against 31 for Flash.

— VentureBeat, November 12
venturebeat.com/technology/google-launches-gemini-4-hydrogen-a-faster-smarter-pricier-flash-successor

BATCH SCHEDULING

Kueue 0.20.6 restores the CRD conversion webhook that 0.20 broke

Without /convert, kube-apiserver failed to build its watch cache for several CRDs. 0.19.13 fixes a quota double charge.

— GitHub release, November 12
github.com/kubernetes-sigs/kueue/releases/tag/v0.20.6

KUBECTL TIP

```
kubectl logs -l name=myLabel
```

Dump logs of pods labelled name=myLabel *(kubernetes.io)*

HPC

Volcano 1.15.4 fixes gang scheduling for DRA-only pods

It also stops a DRF accounting panic and keeps webhook certificates across Helm upgrades. No breaking changes.

— GitHub release, November 12
github.com/volcano-sh/volcano/releases/tag/v1.15.4

BUILD TOOLS

Gradle 9.9.1 and 8.14.7 stop version catalogs from running code

GHSA-3cfg-jpcr-gfp3 (9.3) affects 7.0 to 8.14.6 and 9.0.0 to 9.9.0: catalog values run as code, even during IDE sync.

— GitHub advisory, November 10
github.com/gradle/gradle/security/advisories/GHSA-3cfg-jpcr-gfp3

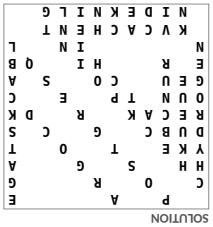
PUZZLE · WORD SEARCH

Ten words from today's two pages are hidden in the grid.

FIND ARGOCD · BACKSTAGE · BATCHING · CHECKPOINT · HYDROGEN · KUEUE · KVCACHE · LINKEDIN · POSTGRESQL · RUNC

M	M	P	P	P	A	E	O	W	R	C	E
C	H	B	O	Y	U	R	A	T	U	C	G
H	H	A	N	S	B	V	G	T	B	O	A
Y	K	E	L	Y	T	N	V	O	L	H	T
D	U	B	C	T	M	G	N	H	C	G	S
R	E	C	A	K	S	T	R	A	A	D	K
O	U	N	C	T	P	N	R	E	L	E	C
G	E	U	O	U	C	O	Y	M	S	R	A
E	W	R	E	R	Y	H	I	I	A	Q	B
N	C	O	V	R	Y	T	I	N	L	A	L
M	K	V	C	A	C	H	E	N	T	N	N
I	N	I	D	E	K	N	I	L	G	U	R

Words run in every direction, also backwards.



SOLUTION