AcornOps

AI helps you ship. Can it keep prod afloat?



Troubleshooting with AcornOps, an open-source AIOps platform you can self-host

Lim Wan Ning, Tan Jin | FOSSASIA 2027

About Us



Platform Engineers @ CSIT, Singapore

















Building software

Developers already build with AI

65.9% of respondents said they use an AI coding assistant or agent at work.

Stack Overflow Developer Survey 2026 · 17,464 answered this question

---

Across three field experiments, developers using GitHub Copilot completed an estimated **26% more pull requests per week.**

Cui et al., Management Science, 2026 · 4,867 developers

Running software

What happens after the code ships?

A deployment can succeed while users are getting errors.

Someone has to spot the failures, find out who's affected and work out what changed.

In DORA's survey, teams using more AI reported higher delivery throughput and less stable releases.

DORA, 2025 software delivery research



Investigating a running service

How AIOps helps when requests fail

Artificial intelligence for IT operations (AIOps) helps teams monitor and troubleshoot running systems.

Spot the spike	An anomaly detector can learn a service's usual error rate and flag an unusual increase.
Find a lead	An assistant can check logs and deployment events from the same time for a clue an engineer can verify.

IBM's AIOps overview · OpenSearch anomaly detection

AI for operations

Which tasks do developers hand to AI?

Share of respondents who said they delegated each task to AI in the previous 30 days

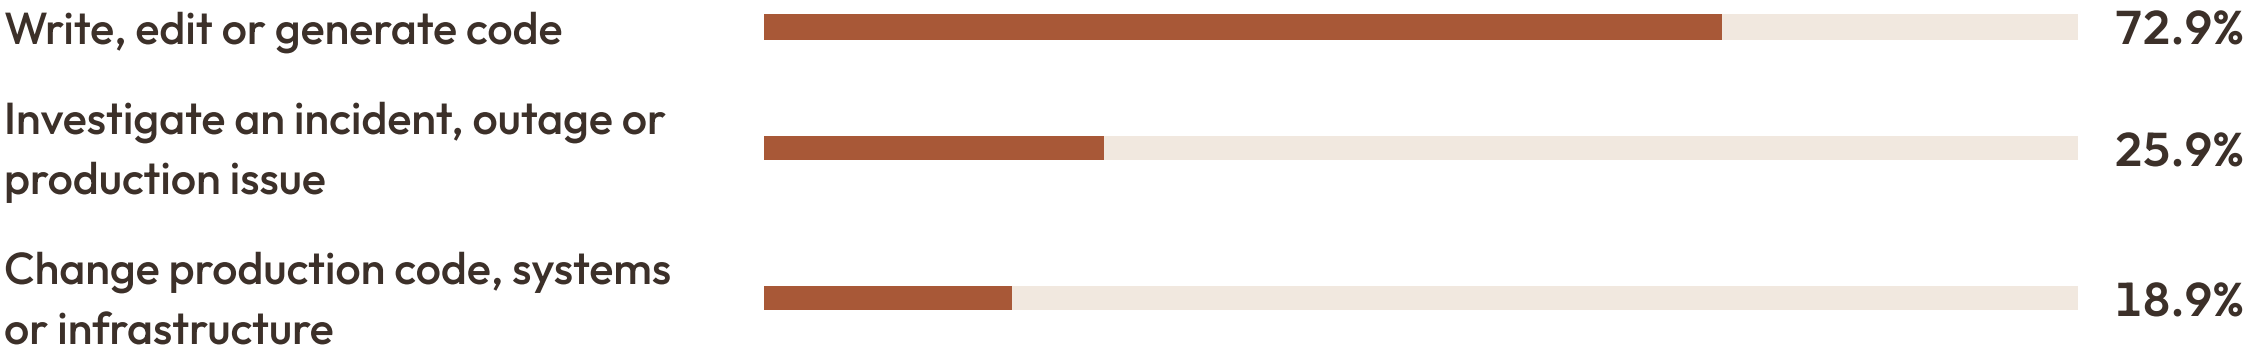

Stack Overflow Developer Survey 2026 · 13,756 answered this optional question · multiple selections allowed

At the whiteboard

A teammate can’t open the app

You’re designing a shared editor.

Your team is debating what should happen when two people edit the same sentence. While AI builds a prototype, a teammate tries the live app.

Your teammate

“I can’t open the app. Is it down for everyone?”

An interrupted afternoon

What does “down” mean here?

One person says the app won’t open. Several failures could look the same on their screen.

Browser	Network	App	Dependency
A script throws an error.	The hostname won’t resolve.	A request handler crashes.	A database call times out.

What would help you tell them apart?

Following the report

How you'd investigate by hand

Try the URL	Ask for the exact address, time and error. Try the same action yourself.
Check the dashboard	Look at traffic and errors around that time. Is one route affected, or are others failing too?
Follow a request	If the request reached the service, check its logs or trace. Where did it first fail?

If a teammate takes over now, how will they know what you've already ruled out?

Building an operations assistant

Why we built AcornOps

The same few engineers knew which checks to run.

When one of them was busy, the next engineer had to piece together what had been tried and why a proposed fix made sense.

Share the checks

Give the assistant the team's instructions and choose where it can look.

Inspect the evidence

See which tools ran and which results you can verify.

Review a change

Require an engineer to approve the exact edit. Then retry the request that failed.

Live demo

Can AcornOps restore the service?

We'll open the failing page on a disposable demo service and follow its workload alert in AcornOps.

We'll check the proposed image change, approve that edit, then try the same URL again.

Open live console ↗

Demo backup

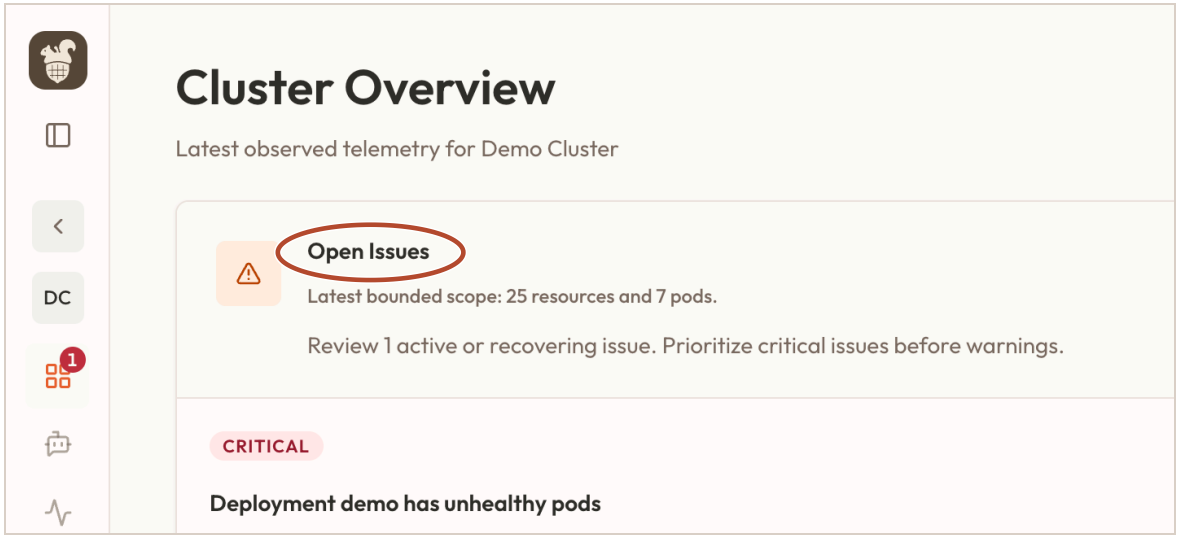

Earlier hosted example · Cluster Overview → Open Issues

Back to the user

Finish with the request that failed

Your teammate wanted to know if the app would open. After the approved image change, we tried the page again.

→ Start with read-only access on a service you know. Before approving a fix, choose the request you'll retry.

Recorded local repair

`GET /` → **HTTP 200**

The page read "Welcome to nginx!"

AcornOps

linktr.ee/acornops

Thank you!

Recorded demo

The container is waiting for its image

ImagePullBackOff.

Kubernetes is retrying the image download. What error did the last attempt return?

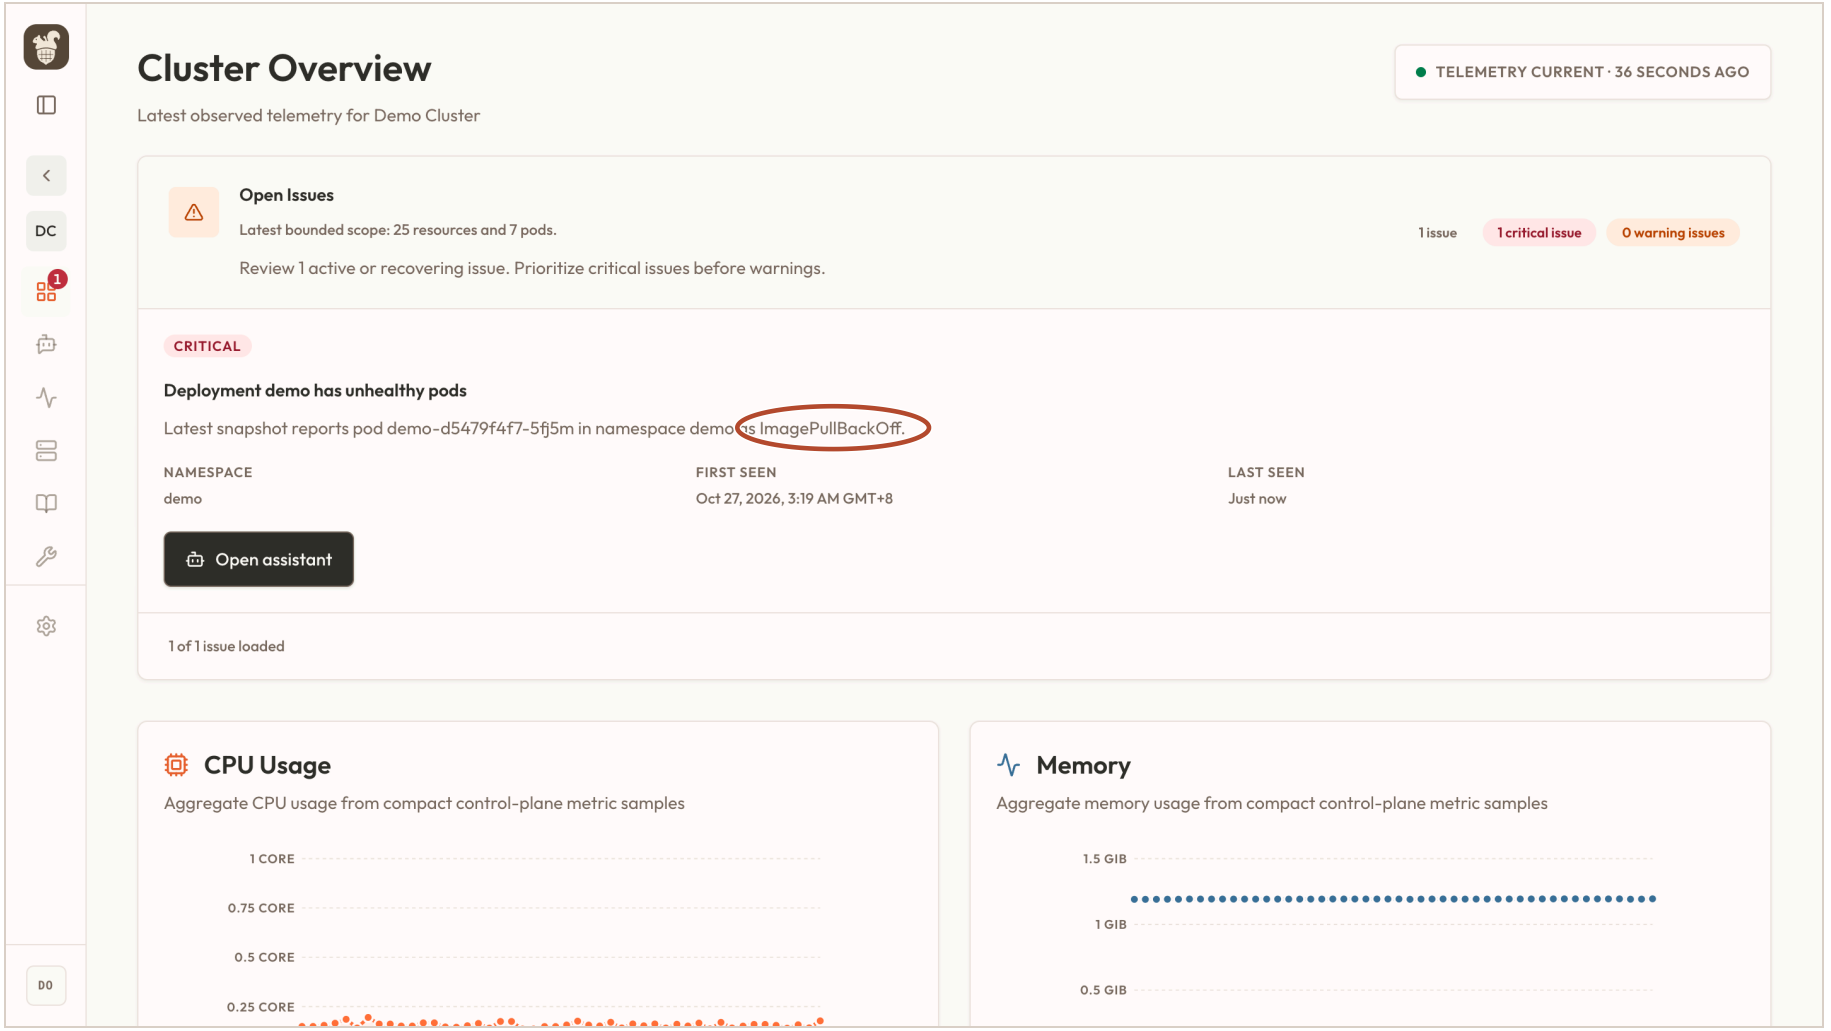

The alert supplies the target

Deployment/demo.

The alert fills in the workload, namespace and symptom.

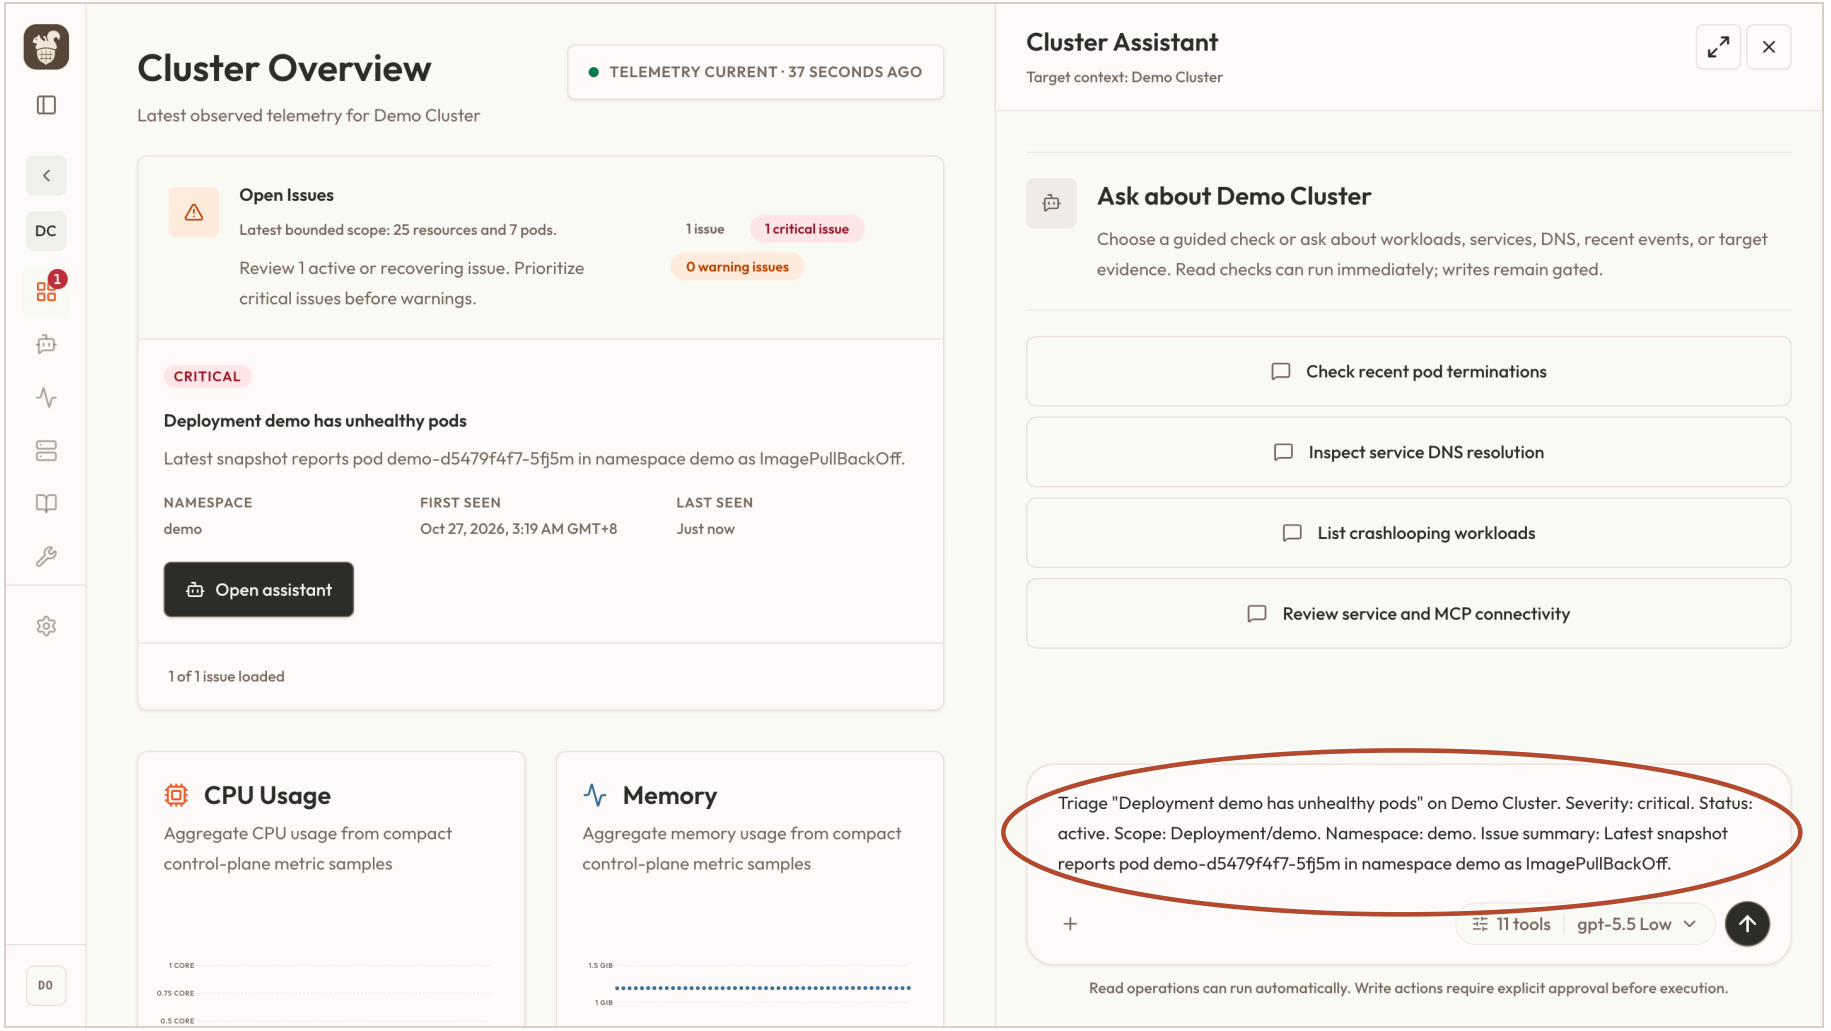

What can this assistant change?

Writes require approval.

This session can request a workload edit. An engineer decides whether it can run.

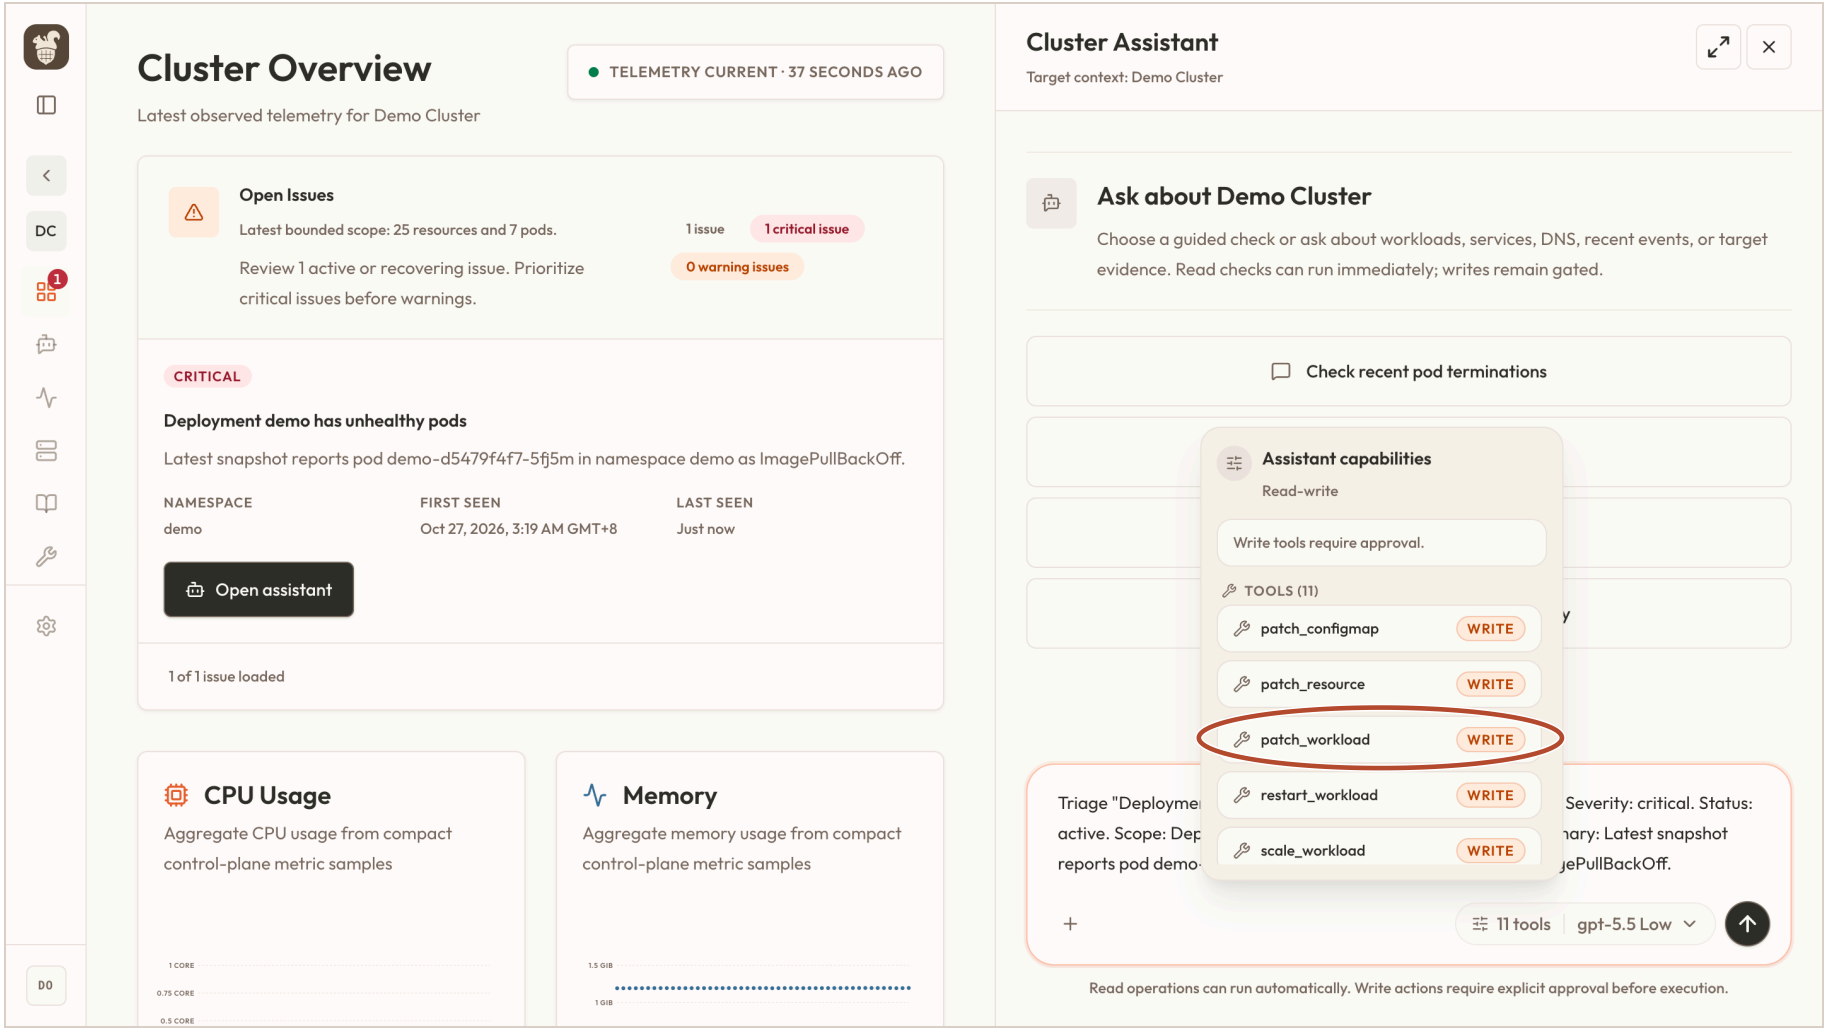

An earlier run reports a missing tag

“Image tag not found.”

The earlier chat names nginx:1.27.4-alpnie. Can we verify that claim?



Can we still check the original read?

“The full result is no longer available.”

We need a fresh read before relying on this diagnosis.

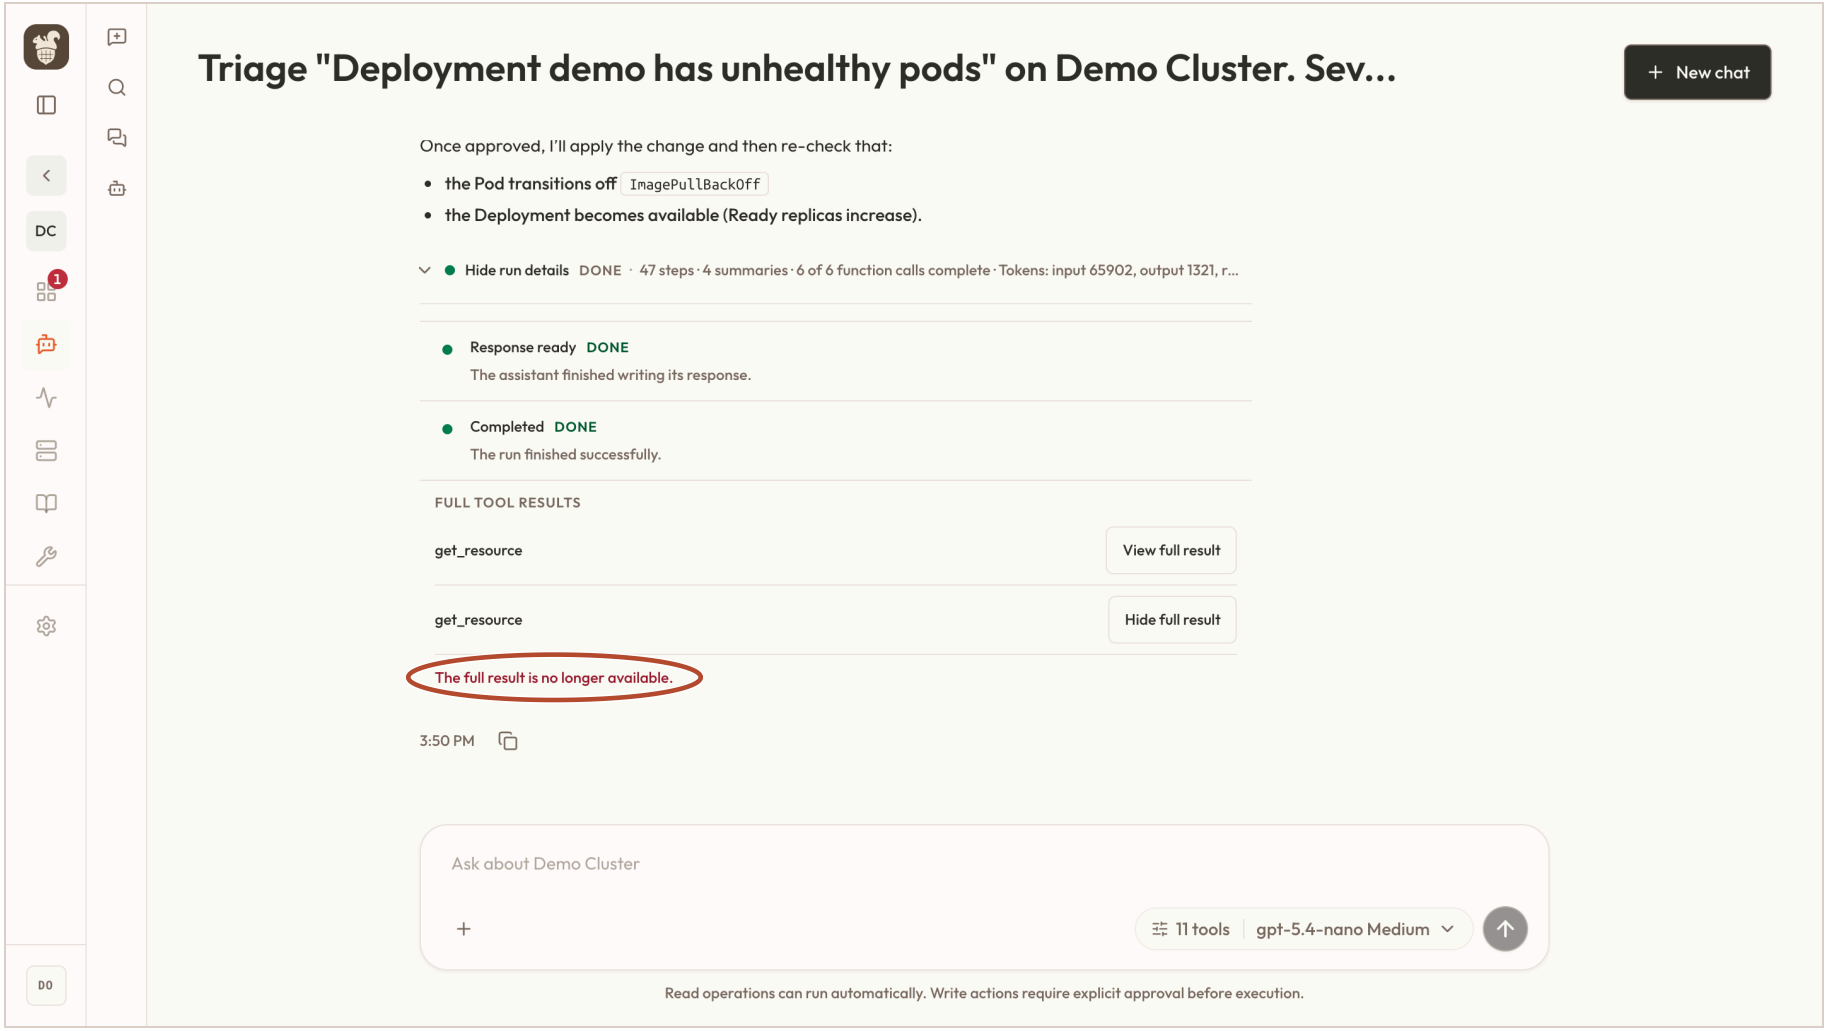

The proposed image change

The edit swaps two letters.

This hosted request was declined. The next step shows the same correction approved in a separate local rehearsal.

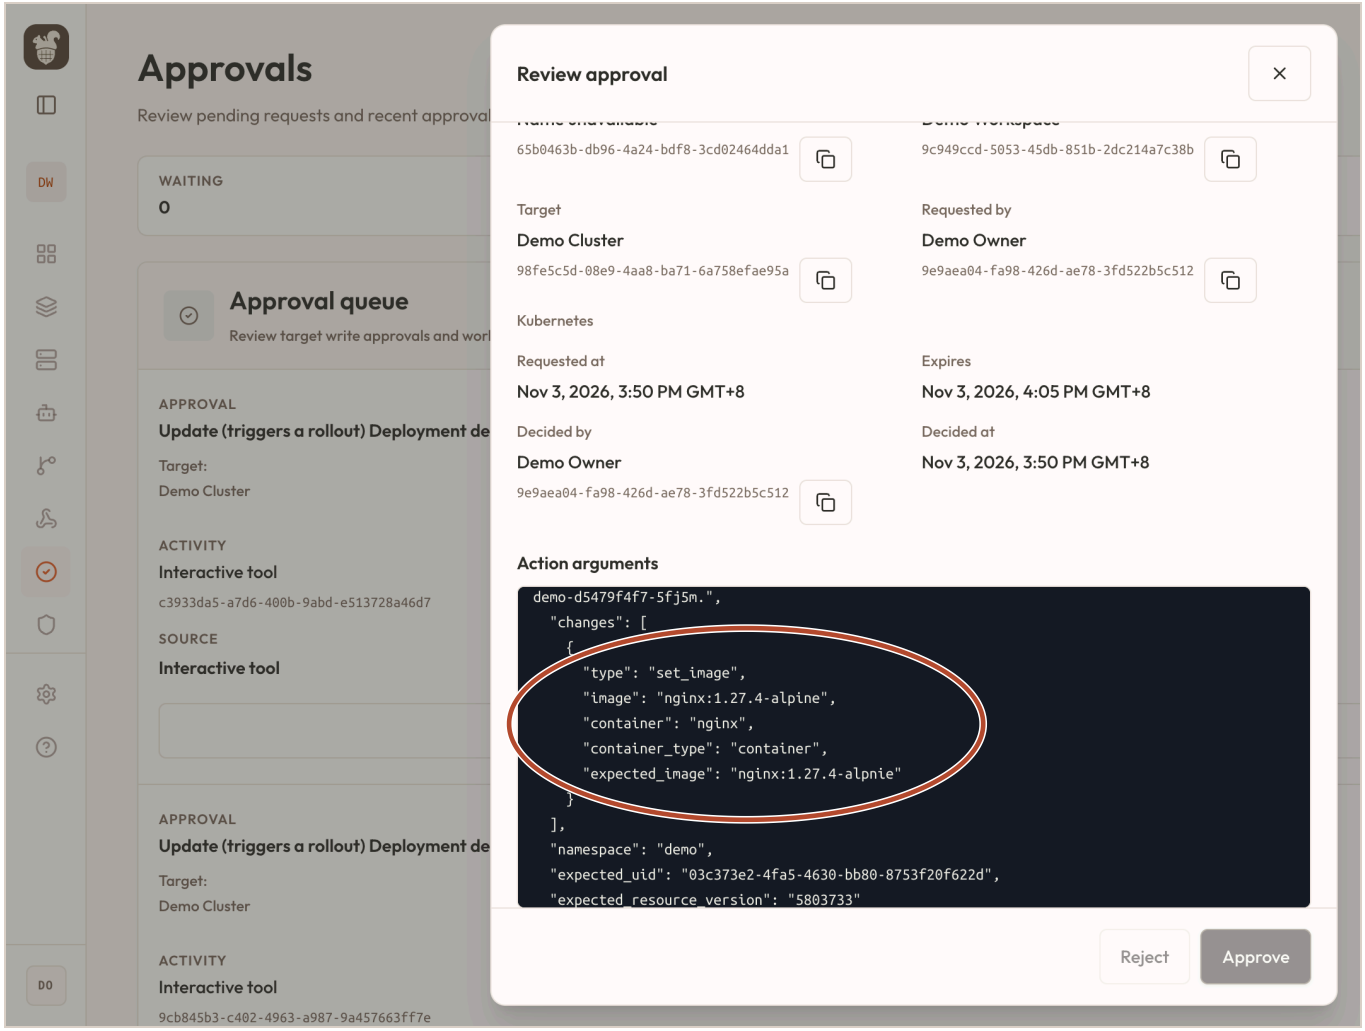

Current tag

`1.27.4-alpnie`

Proposed tag

`1.27.4-alpine`

The service returned the expected page

Before the edit	Kubernetes couldn’t pull `nginx:1.27.4-alpnie`, leaving `demo/web` with 0/1 replicas ready.
Approved change	AgentK applied `nginx:1.27.4-alpine`. A fresh read confirmed the image and 1/1 replicas ready.
Service response	`GET /` through the Service returned HTTP 200. The page read “Welcome to nginx!”