


OpenSRE: An Open-Source Agent That Works Out Why Production Broke

Kabui, Charles

2026-10-21

[Read at ToKnow.ai](#)

A dark blue banner for the OpenSRE project. The title 'OpenSRE: An Open-Source Agent That Works Out Why Production Broke' is in white and light blue. Below it, a subtitle reads 'A command-line AI agent for on-call engineers'. The date 'October 21, 2026' is shown. Three statistics are highlighted in colored boxes: '60+' tools it can query (red), '3 of 5' checks failed, independent test (yellow), and '0.1' version, a public alpha (blue). The ToKnow.ai logo is in the bottom right.

OpenSRE: An Open-Source Agent That Works Out Why Production Broke

A command-line AI agent for on-call engineers

October 21, 2026

60+ tools it can query	3 of 5 checks failed, independent test	0.1 version, a public alpha
----------------------------------	--	---------------------------------------

ToKnow.ai

OpenSRE is an open-source command-line agent that tells on-call engineers why a production system broke. Point it at the tools you already run, about 60 of them from Grafana and Datadog to Kubernetes, AWS and PagerDuty, and when an alert fires it reads logs, metrics, traces and recent deploys, tests a few explanations, then posts its answer to Slack. It runs on a laptop or your own server, accepts Claude, GPT, Gemini or a local Ollama model, and

masks pod, cluster and account IDs before anything reaches an outside model. It is at version 0.1, a self-described public alpha, and ships the training rig alongside the agent because the team argues incident response has no shared benchmark yet.

One2N, an SRE consultancy, spent three months running that build against a small Kubernetes stack, killing Postgres, Redis and the backend one at a time and then all at once. It marked five abilities as pass (+1), partial (0) or fail (-1). OpenSRE failed three, partly passed two and totalled minus three. It blamed a rollback of a Kubernetes release that never happened. It could not read the alert format from the team's monitoring tool without a hand-written conversion step, and it treated an alert storm as separate incidents. With no runbooks or past incidents to draw on, every page started cold. The verdict was defer, not reject.

Since then the project has added long-term memory, runbook-aware reasoning and lessons carried between sessions, which is the area where One2N scored it a fail.

Read More: [SecRespond: AI Agents Spot Security Alerts but Miss Silent Intrusions](#)

Sources:

- [OpenSRE repository and README](#)
- [How to evaluate AI SRE agents for production \(One2N\)](#)
- [The One2N test rig and reproduction kit](#)
- [OpenSRE feature list and documentation](#)
- [Reversible identifier masking docs](#)

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)