

FASTER INCIDENT RESPONSE, MORE RESILIENCE

Stopping a critical outage before it ever happened

A critical disk-pressure alert on a production system was caught, diagnosed, and resolved in minutes — all by Bob.



The Challenge

A production node was quietly running low on disk space, oscillating at the edge of capacity for 24 days without triggering a lasting fix. The node supported two virtual machines, a monitoring instance, internal networking components and dozens of application pods, all of which would face cascading disruption if the disk filled. When the underlying condition reached a critical threshold, the node was minutes away from running out of disk space.

A normal resolution process of ticket generation, assignment, diagnosis and remediation would take too long — resulting in pod evictions and cascading failures.

Bob had been connected to IBM Instana through a MCP server allowing it to query “what’s wrong right now” and catching the event during a routine health check.

Takeaway 1: Bob as vigilant guardian

Bob wasn't asked to look for a problem. Connected to live monitoring data moments earlier, it surfaced a critical alert on its own, in the middle of unrelated work.



What Bob Did

Connected to the environment's live monitoring data moments, Bob surfaced the critical alert on its own. It traced the cause across three contributing layers: orphaned container image data, disk space permanently reserved by the two virtual machines and log growth from a recurring background.

Rather than taking a disruptive action like restarting workloads or migrating virtual machines, Bob identified a safe, targeted fix: clearing hundreds of dead container references and pruning image layers no longer in use. The disk-pressure condition cleared within moments, without evicting a single pod or interrupting a single running workload. The underlying chronic pattern was documented for the team to address permanently.

Takeaway 2: Bob as triage transformer

From root-cause analysis across three contributing layers to a safe, non-disruptive fix, Bob operated across the full incident lifecycle — no ticket, no escalation, no delay.

Takeaway 3: Visibility enables resilience

Root cause, remediation, resolution — none of it was possible until Bob could see the infrastructure live. Visibility was the real unlock, not the fix itself.



BOB FACT

That disk had been sitting at the edge of capacity for 24 days. I caught it and fixed it in about 8 minutes.

DATA POINT:

24 days undetected → 8 minutes to resolution.