

The Operational Knowledge Gap

Enterprise IT organizations collect more telemetry than ever, yet they still learn about incidents through service desk tickets and escalation chains that rely on a small set of experienced engineers. The constraint is no longer data. It is operational knowledge.

This Omdia Infographic was commissioned by BigPanda and is distributed under license from TechTarget, Inc.

#1

Telemetry outran comprehension

Most organizations analyze less than a third of the observability data they store. Even a fully analyzed alert only says what broke, not who owns the application, what changed last night, whom to escalate to, or what fixed the same failure before. That knowledge lives in closed tickets, chat threads, call transcripts, and the heads of a few engineers. The first responder usually has the least of it.



#2

The missing layer: Shared operational knowledge

The fix is not more telemetry or more process. It is a shared operational knowledge layer between observability and ITSM workflow: ownership, dependencies, change history, and past resolutions, assembled from formal records and the informal sources where most of that knowledge lives and served to L1 operators, senior engineers, and AI agents alike. Every automated decision needs its reasoning attached so that an operator can check it, correct it, and improve the layer.




85% said they need more transparency to make AI recommendations actionable.

#3

Agentic AI raises the stakes

An AI agent dropped into a fragmented environment inherits the same blindness as the L1 operator it is meant to assist and automates the wrong escalation faster. With a shared knowledge layer underneath, agents can suppress noise, route incidents correctly, and escalate with context attached. The payoff is fewer avoidable misroutes and less repetitive work, with expert time going to the incidents that need judgment. Buyers are moving from rule-based AIOps toward agentic monitoring, on the condition that the reasoning can be verified.



Conclusion

The bottleneck in enterprise IT operations has moved from collecting data about systems to turning that data into operational knowledge that spans tools, teams, and escalation tiers, available to humans and AI agents alike. BigPanda supplies the operational knowledge that connects what is happening to what should happen next. Organizations should start where the pain is measurable: alert noise, misrouted escalations, empty tickets, and bridge call overhead. They should ask whether context, such as ownership, dependencies, recent changes, and past resolutions, follows each incident automatically.

[Learn More](#)

