

SEPTEMBER 2026

From Telemetry to Operational Knowledge: Why Agentic IT Operations Needs a Context Layer

Torsten Volk, Principal Analyst

Abstract: Enterprise IT organizations collect more telemetry than they can use. Omdia research finds that 61% of organizations analyze 30% or less of the observability data they already collect and store.¹ Meanwhile, 81% of organizations identified agentic AI as a strategic priority across IT and cloud.² Making agents effective in incident response requires more than better use of telemetry: It needs the operational knowledge (i.e., who owns an application, what changed, whom to escalate to, and what worked last time) that resides in closed tickets, archived chat threads, and the heads of a few experienced engineers. Telemetry alone cannot reliably supply this context.

This Omdia Showcase examines why operational knowledge has become the binding constraint in enterprise incident response and how the BigPanda agentic IT operations platform addresses it.

The problem: Telemetry outran comprehension

Enterprise IT organizations collect more telemetry than ever, yet they still learn about incidents from customer complaints and service tickets and resolve them through escalation chains that rely on a small set of experienced engineers.

Omdia research finds that 61% of organizations analyze only 30% or less of the observability data they collect and store.³ But closing that analysis gap is only part of the solution.

Even a perfectly analyzed alert, built on complete metrics, events, logs, and traces, only tells an operator what broke. It does not say who owns the application, what changed last night, whom to escalate to, or what resolved the same failure eight months ago. That knowledge exists, but it lives in closed tickets, chat threads, bridge call recordings, and institutional memory.

This operational knowledge gap gets expensive because it sits on top of an already fragmented tool environment. 76% of organizations reported that more than a fifth of their incidents require correlating data across multiple tools, and one-third of organizations run 11 or more monitoring tools.⁴

The operational knowledge that completes the picture of an incident is not telemetry in any tool, and no single team holds it either. Enterprises built ITIL tiering, network operations centers, and L1/L2/L3 escalation

¹ Source: Omdia Research Report, [Unlocking Observability in the Era of AI and Application Modernization](#), March 2026.

² Source: Omdia Research Report, [IT Modernization Report: Cloud, Applications, and Infrastructure Amid the AI Revolution](#), August 2026.

³ Source: Omdia Research Report, [Unlocking Observability in the Era of AI and Application Modernization](#), March 2026.

⁴ Ibid.

chains precisely to protect application engineers from operational interruptions. The result is a structural mismatch. The first responder to an incident is usually the person with the least context about the affected application, while the only people permitted to touch the code are the most expensive to interrupt.

Agentic AI raises the stakes on this mismatch instead of resolving it. 81% of organizations called agentic AI a strategic priority across IT and cloud, and 80% said ensuring positive ROI for their AI initiatives is a top concern.⁵ An AI agent dropped into a fragmented environment inherits the same blindness as the L1 operator it is meant to assist. Without a shared operational context, agents automate the wrong escalation faster.

Consider a retail payment authorization service that degrades at 2 am. Alerts fire in the APM tool, the network monitor, and the cloud provider console. The L1 operator in a centralized NOC sees three disconnected fragments of one incident. The application team that could interpret them is asleep in another time zone, and the runbook covers a version of the service retired last year. The actual fix, a connection pool restart, sits undocumented in an archived Slack thread from a similar incident eight months earlier. The outage runs for four hours. The knowledge to resolve it in fifteen minutes existed the entire time, but the organization had no way to find it.

An AI agent dropped into a fragmented environment inherits the same blindness as the L1 operator it is meant to assist. Without a shared operational context, agents automate the wrong escalation faster.

Figure 1. Enterprise IT incident resolution challenges



Source: Omdia

Why the obvious fixes fall short

More telemetry data does not close the gap and neither does consolidating observability into one platform. Both are worth doing, and both hit the same ceiling. Complete, unified metrics, events, logs, and traces still describe only system state. Ownership, escalation paths, change history, and prior resolutions never enter the telemetry stream, so a single perfect telemetry platform still leaves the operator uninformed about what to do next.

⁵ Source: Omdia Research Report, *IT Modernization Report: Cloud, Applications, and Infrastructure Amid the AI Revolution*, August 2026.
©2026 TechTarget, Inc. All Rights Reserved.

ITSM workflow alone does not solve this problem either. Ticketing and workflow platforms define the process an incident should follow, but they do not know when to trigger it or what context to attach. Tickets arrive empty, and humans fill them by asking around.

More headcount does not close it either. Tribal knowledge does not scale by hiring, and it walks out the door when people leave. Every new operator restarts the same learning curve against the same insufficiently documented environment.

Classic AIOps, understood as rule-based event correlation, is losing buyer confidence. When Omdia asked about planned network management investments, agentic AI-enabled monitoring and management ranked first, cited by 34% of organizations, while AIOps trailed at 19%.⁶ Buyers are asking for systems that understand their environment.

Ownership, escalation paths, change history, and prior resolutions never enter the telemetry stream, so a single perfect telemetry platform still leaves the operator uninformed about what to do next.

What any credible solution must provide


The pattern across the failed fixes points to a missing layer between observability and workflow, one that carries operational knowledge rather than more telemetry or more process. Three requirements define whether a product genuinely occupies that layer or just relabels an adjacent one:

- **A shared operational knowledge layer spanning tools, teams, and tiers.** The knowledge needed to resolve an incident is scattered across observability platforms, the CMDB, change records, service desk tickets, and collaboration channels. A credible solution has to assemble it in one place and serve it to every consumer, from the L1 operator to the senior engineer to an AI agent.
- **Tolerance for messy, incomplete, and informal knowledge.** Any approach that requires a curated CMDB or complete documentation as a prerequisite disqualifies itself as a credible solution because no large enterprise has either of these. The most valuable operational knowledge is unstructured: the bridge call where someone explains the workaround, the Slack chat where the real dependency surfaces, the closed ticket with the fix that never made the runbook. The knowledge layer must capture and improve as a byproduct of daily operations rather than as a documentation project.
- **Agent-ready context with reasoning that humans can verify.** 85% of organizations said they need more transparency to make AI recommendations actionable.⁷ Humans and AI agents must draw on the same governed source of knowledge, and every automated decision needs its reasoning attached so an operator can validate it, correct it, and, by doing so, improve the layer for the next incident.

⁶ Ibid.

⁷ Source: Omdia Research Report, [Unlocking Observability in the Era of AI and Application Modernization](#), March 2026.

Figure 2. Unified knowledge for incident resolution



Source: Omdia

How the BigPanda agentic IT operations platform fits in

BigPanda positions its agentic IT operations platform against exactly this operational knowledge gap. Its platform consists of four product areas: AI Incident Prevention, AI Detection and Response, L1 Agent, and AI Incident Assistant, all built on a common data foundation called the BigPanda IT Knowledge Graph. BigPanda describes two of the four, AI Detection and Response and L1 Agent, as working together as one frontline detection, triage, and response capability, and this paper treats them that way. BigPanda does not position the platform as a replacement for observability tools or ITSM platforms and says that most of its customers run it alongside LogicMonitor, Dynatrace, ServiceNow, or comparable products. The positioning claim is specific: Observability tells teams what is happening, ITSM defines the workflow to follow, and BigPanda occupies the layer between them, supplying the operational knowledge that connects the two.

- **The IT Knowledge Graph** is the mechanism behind the first two requirements of a credible agentic IT operations solution, a shared knowledge layer and tolerance for informal knowledge. It ingests and connects structured operational data, including topology, CMDB entries, and change records, together with service desk tickets, runbooks, and architecture diagrams. It also captures the knowledge that usually stays tribal: chat conversations, bridge call transcripts, and incident postmortems. The design point that matters most, in Omdia's reading, is that the graph is not a one-time integration of existing sources. BigPanda describes it as learning continuously from what it observes in operation. Each incident that is investigated, resolved, and documented adds to the graph: which signals belonged to the same incident, which team actually owned the fix, which change was behind the failure, and what resolved it. Operator corrections feed back the same way. BigPanda states that the graph is designed to work from incomplete and informal inputs rather than requiring a curated CMDB or complete documentation as prerequisites. If that holds in practice, the context available to humans and agents improves with every incident instead of decaying the way written documentation does.

The intended effect is that an incident reaches a human with the context already attached: the owner, the dependencies, last night's changes, and the fix that worked before. Omdia's evaluation question is how much of this learning happens without human effort. Buyers should probe how much knowledge capture occurs automatically versus how much requires operator curation, how quickly the graph retires stale or

incorrect knowledge, and whether they can see what the graph learned from a specific incident after the fact.

- **AI Detection and Response and L1 Agent** together form the frontline detection, triage, and response capability that consumes that knowledge where incidents first arrive. In BigPanda's description, AI Detection and Response correlates signals from observability tools, service desk activity, and external dependencies such as provider outages, then runs an automated investigation, which BigPanda calls AI Triage, that attaches a summary, similar past incidents, and the change most likely responsible. L1 Agent picks up where that investigation ends and acts on it, suppressing noise and false positives, prioritizing and routing incidents, resolving a subset autonomously, and escalating the rest with the investigation attached. BigPanda states that each automated action carries its reasoning so that operators can review a decision rather than accept it unexamined, which maps directly to the verifiability requirement noted above. The outcome worth measuring is not fewer people at L1 but fewer avoidable misroutes and less repetitive investigation so that L1 operators spend their time on the incidents that need judgment and L3 engineers receive escalations with context rather than empty tickets.
- **AI Incident Assistant**, which includes Biggy AI as the interactive agent inside Slack, Microsoft Teams, and the BigPanda web application, serves the escalation tiers. L2 and L3 engineers, SREs, and incident commanders investigate in natural language, pull in responders identified from historical incident context, and build workflow automations without code. AI Incident Prevention extends the same knowledge foundation into the future by scoring the risk of proposed changes before rollout.

BigPanda cites a public-sector example: the California Department of Motor Vehicles, where incident response times reportedly fell from hours to 15 minutes. On economics, BigPanda argues that enterprises collectively spend \$250bn annually on manual, human-driven IT operations, and a BigPanda-commissioned business value study of 23 enterprise customers reports a median ROI of 430% with payback in under a year. Set against the 80% of organizations that named ensuring a positive ROI for AI initiatives a top concern,⁸ these are the claims a buyer will want to test against their own escalation volumes during evaluation.

What buyers should still verify

The trust gap between stated confidence and operational behavior deserves attention, as 86% of organizations said they are comfortable letting AI autonomously make low-risk configuration changes.⁹ Nearly all organizations (95%) said they were confident or very confident in their ability to manage AI agent risks.¹⁰ Yet 46% still said they always or often override AI recommendations.¹¹ Stated confidence and actual delegation are different things.

The benchmark buyers apply deserves the same scrutiny. Few organizations measure the accuracy of their own human triage, even though the structure described in this paper implies it is weak: The first responder is the operator with the least context about the affected application, working from empty tickets across disconnected consoles. Buyers should quantify that baseline, meaning current misroute, reopen, and false-escalation rates at L1, and judge agent accuracy against it instead of against an assumed standard of

⁸ Source: Omdia Research Report, [IT Modernization Report: Cloud, Applications, and Infrastructure Amid the AI Revolution](#), August 2026.

⁹ Source: Omdia Research Report, [Unlocking Observability in the Era of AI and Application Modernization](#), March 2026.

¹⁰ Source: Omdia Complete Survey Results, [AI Agents: The Game-changing Generative AI Use Case](#), August 2025.

¹¹ Source: Omdia Research Report, [Transforming Observability and Monitoring Through AI](#), April 2025.

flawless human judgment. Autonomy should still expand gradually, starting with suppression and routing, with remediation held to a higher evidence bar.

Two further points belong in any honest evaluation. First, the ROI figures in this space, BigPanda's included, are vendor-commissioned and require independent verification. Second, BigPanda partners with ServiceNow to advance the path to autonomy. In the stated division of labor, ServiceNow provides workflow governance and change control, while BigPanda supplies the operational knowledge and agent context that enable workflows to execute autonomously. The two products still meet at the workflow boundary, so buyers with a deep ServiceNow footprint should map where each one owns the workflow before assuming clean complementarity.

Conclusion

The bottleneck in enterprise IT operations has shifted from collecting system data to turning that data into operational knowledge that spans tools, teams, and escalation tiers and is available to humans and AI agents alike. The BigPanda agentic IT operations platform is a deliberate attempt to address that layer. In BigPanda's stated design, the IT Knowledge Graph assembles context from structured and unstructured sources and learns from each incident it observes, and the platform's agents consume that context across the incident lifecycle.

Enterprises with centralized NOC or outsourced L1 GSI structures, double-digit tool counts, and measurable escalation costs should evaluate first, beginning where the pain is quantifiable: misrouted escalations, empty tickets, and bridge call overhead. Organizations with a single consolidated observability stack and developer-owned on-call rotations will see less benefit because their first responder already has code-level context. Organizations unwilling to grant a platform access to tickets, chat, and change data across silos should also wait, since a knowledge layer starved of sources cannot deliver on its premise.

From Omdia's perspective, the operational knowledge layer is emerging as a distinct category between observability and ITSM. Within the next two years, most vendors in both adjacent categories will claim it, and every one of them will demonstrate an agent that closes an incident. That demonstration proves little on its own. A better test is to take a real incident from last quarter and ask two questions. First, did the agent's decision draw on operational knowledge? Did it know who owned the service, what changed, and what fixed it last time, or did it act on the alert alone? Second, can the reasoning be verified? Can an operator see why the agent suppressed, routed, or escalated and correct it so the next incident goes better? A platform that passes both tests has a knowledge layer. A platform that fails either one has an agent acting on alerts, whatever the vendor calls it.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the “Omdia Materials”) are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together “Informa TechTarget”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

Get in touch: www.omdia.com askananalyst@omdia.com

