VIEWPOINT



# WHY AIOPS STALLS: THE NEED FOR A TRUSTED OPERATIONAL EVIDENCE LAYER

Enterprises can unlock the full value of AIOps by creating a trusted operational evidence layer that connects telemetry, business context, and governance.

Infosys
Navigate your next

# Executive Summary

Organizations are investing heavily in AI-powered IT Operations with the expectation of reducing outages, accelerating root-cause identification, improving operational efficiency, and eventually enabling autonomous operations. Investment momentum is strong, yet enterprise-scale adoption remains limited. A 2025 survey found that organizations nearly doubled AI investments and reported positive AIOps outcomes, but only a small percentage of initiatives progressed to full deployment. [1]

The reasons are becoming increasingly clear:

- AI is being deployed on operational environments that remain fragmented, inconsistent, and poorly contextualized.
- While organizations focus on AI platforms, many overlook the quality of the operational evidence on which those platforms base their decisions.
- AI can only be as reliable as the telemetry, service context, and governance behind it.

This point of view captures why the primary barrier to enterprise-scale AIOps is not AI capability but the absence of a **Trusted Operational Evidence Layer** that provides AI with complete, contextualized, and governed operational intelligence.



External Document © 2026 Infosys Limited



# The AI Operations Paradox

Organizations have never had more operational data. At the same time, they still struggle to understand end-to-end service behavior.

Modern enterprises operate across cloud and on-premise infrastructure, SaaS platforms, containers, APIs, edge environments, and increasingly, AI-enabled workloads. Every domain generates telemetry, yet that telemetry is often distributed across multiple tools, teams, and operating models.

Research indicates that organizations use an average of 13 observability tools from 9 vendors while actively pursuing consolidation. [2]

The result is what can be described as operational visibility debt: the widening gap between what the business needs to understand and what operational systems can explain.

This debt appears as:

- Missing telemetry and monitoring blind spots
- Duplicate and conflicting alerts
- Incomplete service dependency maps
- Stale configuration management database (CMDB) records
- Inconsistent ownership models
- Fragmented dashboards
- Weak linkage between incidents, changes, and business impact

When AI is layered onto this environment, it inherits and can further amplify the same limitations.

Organizations often expect AI to identify root causes, predict failures, and automate remediation. However, AI ultimately reasons over the same fragmented operational evidence that already constrain human operators.

The outcome is predictable: **AI ambition rises faster than operational readiness.**


External Document © 2026 Infosys Limited

# Why AIOps Fails

In enterprise environments, AIOps initiatives often struggle to scale for four interrelated reasons:

## Fragmented Evidence

Logs, metrics, traces, events, tickets, topology data, and change records frequently reside across multiple systems with inconsistent naming conventions, service definitions and ownership models. AI cannot reliably correlate signals or reconstruct service behavior when its evidence is fragmented at source.[3]

## Poor Data Quality

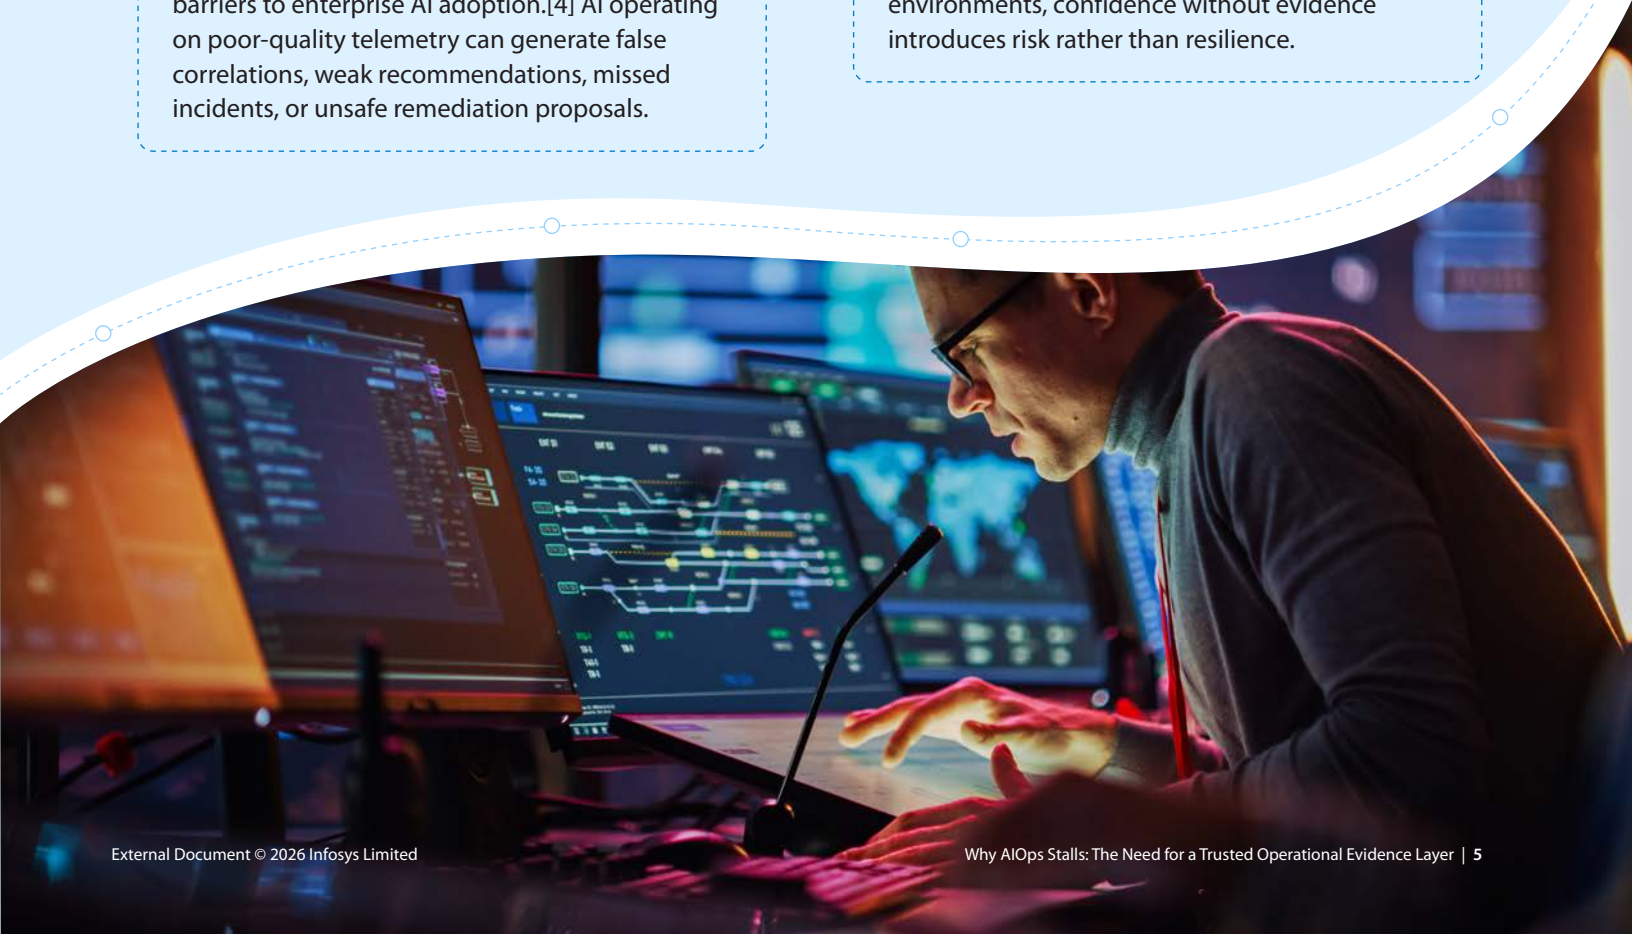
Even when telemetry is available, it may be incomplete, duplicated, delayed, inconsistently labeled, or weakly governed. Research identifies data-quality and readiness gaps as significant barriers to enterprise AI adoption.[4] AI operating on poor-quality telemetry can generate false correlations, weak recommendations, missed incidents, or unsafe remediation proposals.

## Lack of Business Context

AI may identify that a server, application, or container is degraded. Without service and business context, however, it may not know which customer journey is affected, which business outcome is at risk, or whether technically valid remediation is operationally appropriate. Without business context, AI can identify symptoms but struggle to understand consequences.

## Automation Before Trust

Organizations may attempt to automate operational actions before establishing confidence in the underlying evidence. This creates the appearance of intelligence without the evidentiary foundation and controls required for trustworthy action. Generative AI can produce plausible summaries and recommendations, but language models can confidently produce incorrect conclusions when uncertainty is not adequately represented. In operations environments, confidence without evidence introduces risk rather than resilience.



External Document © 2026 Infosys Limited

# The Trusted Operational Evidence Layer

The path forward requires more than another monitoring tool or AI platform. Enterprises need a trusted layer that brings together telemetry, topology, ownership, change data, incidents, business services, and operational knowledge, into a single evidence model so that AI can reason over consistent and contextualized evidence.

Its purpose is simple: **to convert operational data into trusted operational intelligence.**

A Trusted Operational Evidence Layer consists of three foundational capabilities:

## Unified Telemetry

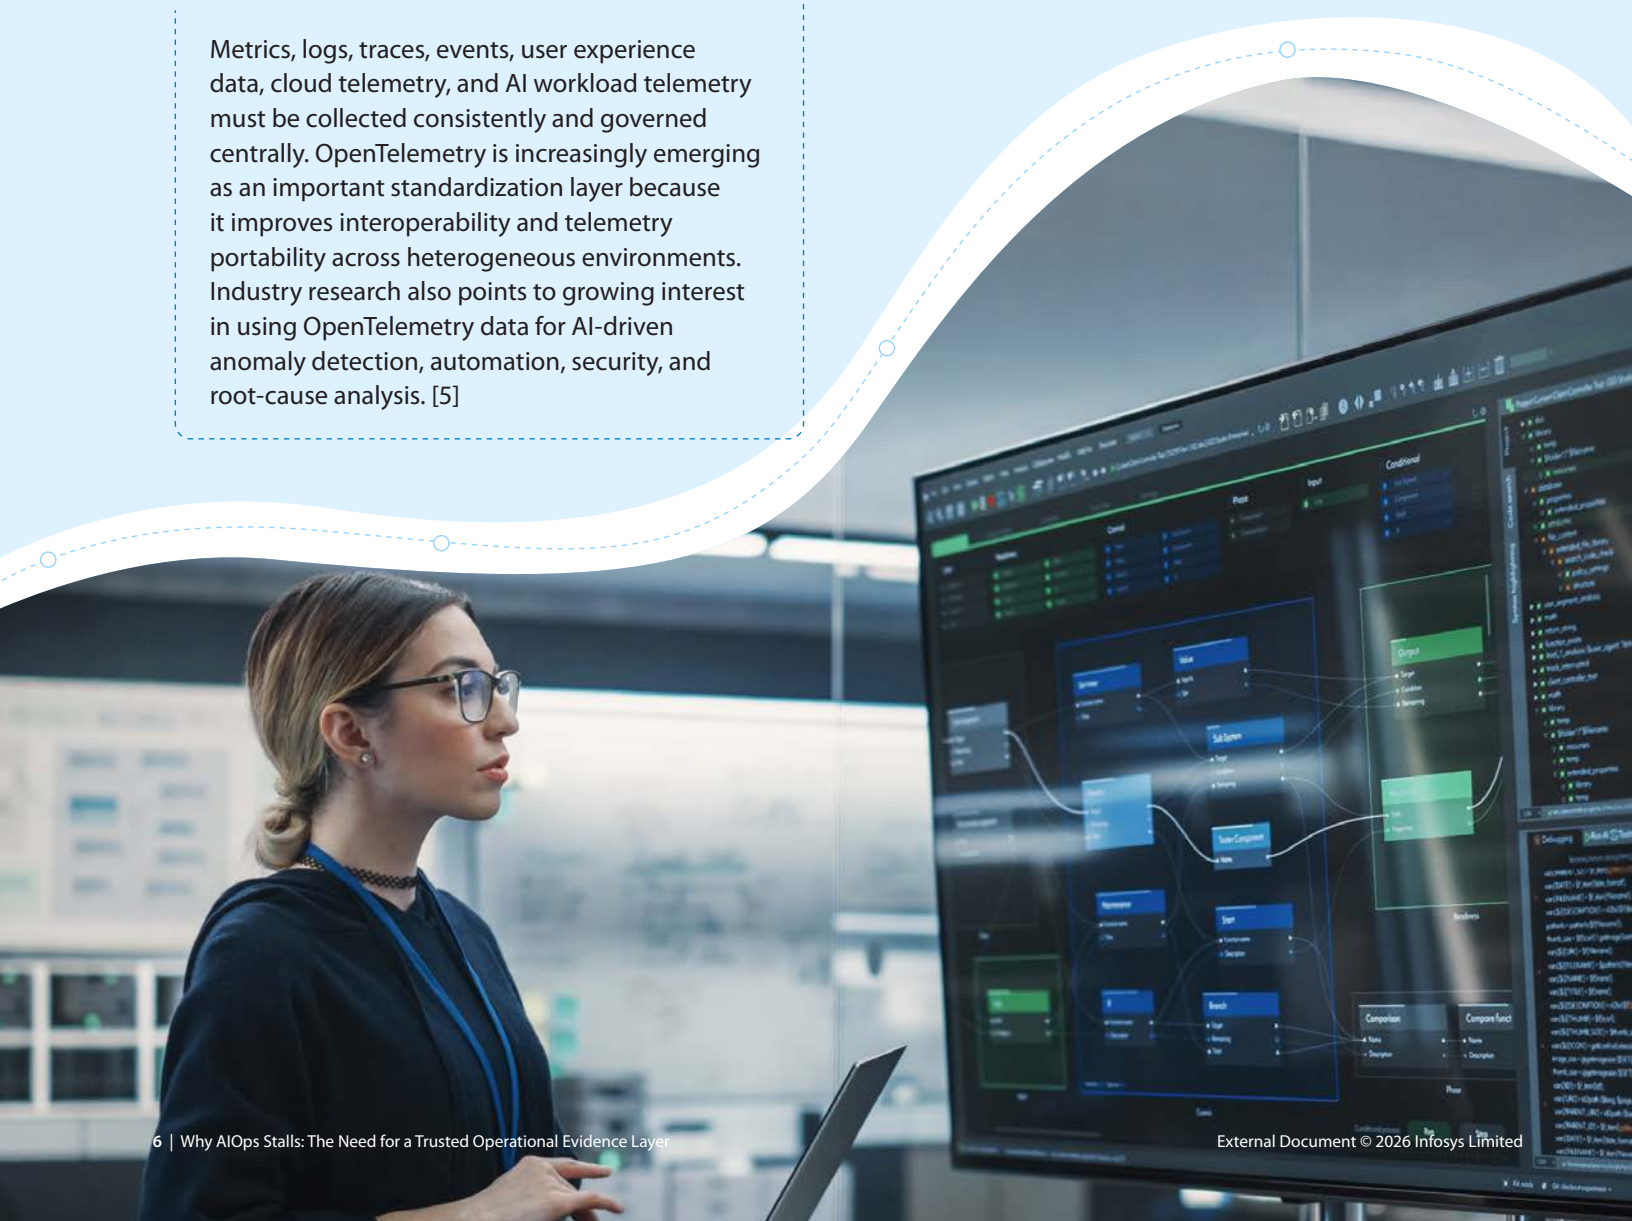
Metrics, logs, traces, events, user experience data, cloud telemetry, and AI workload telemetry must be collected consistently and governed centrally. OpenTelemetry is increasingly emerging as an important standardization layer because it improves interoperability and telemetry portability across heterogeneous environments. Industry research also points to growing interest in using OpenTelemetry data for AI-driven anomaly detection, automation, security, and root-cause analysis. [5]

## Service Context

Telemetry alone rarely provides sufficient insight into business impact. To become meaningful for AI-driven operations, operational data must be enriched with:

- Service ownership
- Application dependencies
- Infrastructure relationships
- Change history
- Business criticality
- Customer journey mapping

Together, these dimensions transform raw signals into contextualized operational evidence.




External Document © 2026 Infosys Limited



## Trust and Governance

AI recommendations should always be explainable, evidence-based, and proportionate to the risk of the proposed action.

- Organizations require:
- Confidence thresholds
- Human approval controls
- Auditability
- Rollback mechanisms
- Model monitoring
- Governance policies

These controls help ensure that AI-generated recommendations remain transparent, accountable, and aligned with enterprise risk tolerance.

Trust is not created at the point of automation. It must be embedded into the operational evidence itself, enabling organizations to scale AI adoption with confidence. At Infosys, we believe that observability must evolve from a monitoring discipline into an evidence discipline for AI-powered operations. The future of AIOps will be shaped by the quality of the operational evidence available to them.

While many organizations focus on deploying AI capabilities, the greater challenge lies in creating an operational environment where telemetry, service context, governance, and business insights are unified into a trusted foundation for decision-making.

The organizations that realize the greatest value from AIOps will not necessarily be those deploying the most advanced AI models, but those that establish the strongest foundation of trusted operational evidence on which those models can act.

External Document © 2026 Infosys Limited

# The Leadership Agenda

For CXOs, the implication is clear. AIOps investment should be evaluated not only by model capability or automation potential, but by the quality of the operational evidence supporting it. Four priorities should guide the agenda:

## Reduce Visibility Debt

Identify and eliminate telemetry gaps, ownership ambiguity, duplicate alerts, stale dependency maps, and fragmented operational views. Prioritize gaps based on their effect on critical services and customer journeys.

## Govern Telemetry as a Strategic Asset

Treat operational telemetry with the same rigor applied to enterprise data, including quality standards, ownership, retention policies, access controls, and cost governance.

## Build Business Service Context

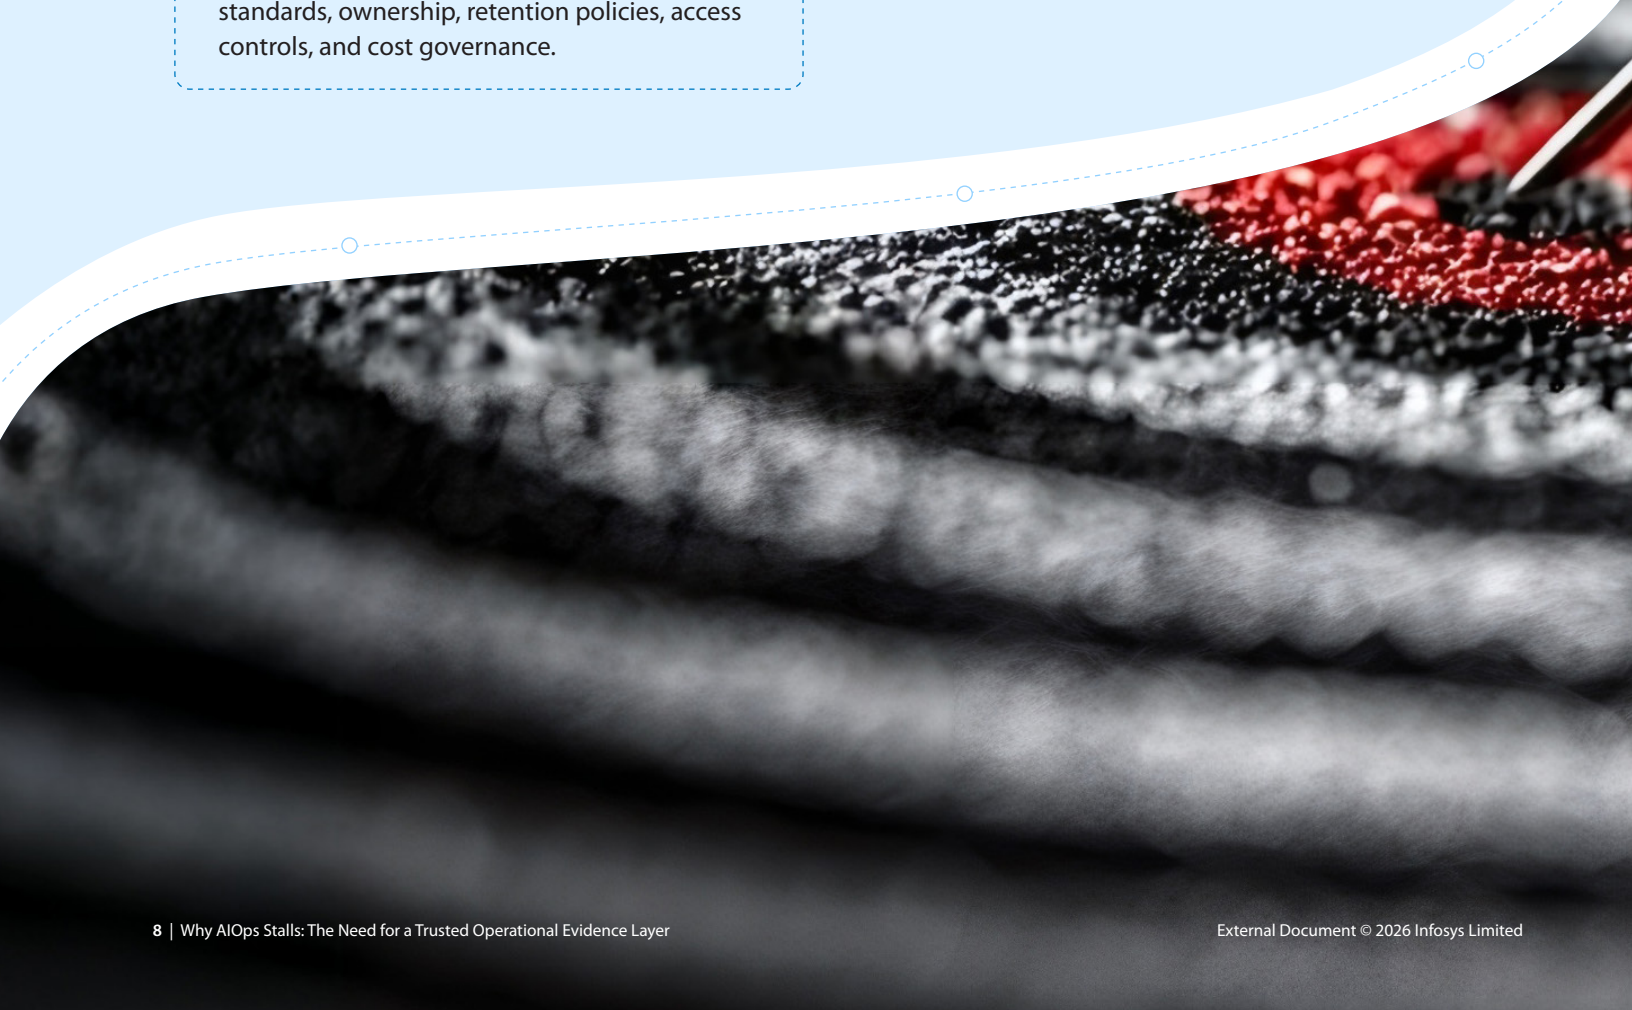
Connect infrastructure and applications signals to services, customer journeys, resilience objectives, and business outcomes.

## Establish an Operational AI Trust Framework

Define the evidence required for AI recommendations, the conditions requiring human approval, the actions eligible for automation, and the controls for auditing decisions and monitoring model quality.

The objective is not autonomous operations.

The objective is trusted operations.





External Document © 2026 Infosys Limited



# Conclusion

AI-powered ITOps will not become trustworthy simply by adding more sophisticated models or more autonomous agents. It will advance when organizations can give AI complete telemetry, business context, governed evidence, and clear limits for action.

Observability is therefore evolving beyond monitoring. It is becoming a critical evidence and control layer for AI-enabled operations.

The critical leadership question is no longer:

Which AIOps platform should we deploy?
It is:

Do we have a trusted operational evidence layer that enables AI to make decisions affecting resilience, customer experience, revenue, and risk?
Until that foundation exists, AI-powered ITOps will remain a promising vision rather than an operational reality.

External Document © 2026 Infosys Limited

## References


1. The Future of IT Operations in the AI Era | Riverbed Global Survey '25

2. The Future of IT Operations in the AI Era | Riverbed Global Survey '25

3. https://www.dynatrace.com/info/whitepapers/idc-from-observability-to-supervised-autonomous-operations/

4. The Future of IT Operations in the AI Era | Riverbed Global Survey '25

5. The Future of IT Operations in the AI Era | Riverbed Global Survey '25

6. Gartner, Magic Quadrant for Observability Platforms (July 2025). Observability platforms are evolving through analytics, cost optimization, and AI observability. [Link: Gartner Magic Quadrant for Observability Platforms]

7. Riverbed, The Future of IT Operations in the AI Era (2025 Global Survey). AI investment, deployment maturity, tooling fragmentation, OpenTelemetry adoption, and operational readiness findings. [Link: The Future of IT Operations in the AI Era | Riverbed]

8. IBM, What is Observability in AIOps? Definition and role of telemetry in anomaly detection, RCA, predictive analytics, and automation. [Link: What is observability in AIOps? | IBM]


## About the Author




**Vivin George**

*Infosys*


AVP and Senior Principal Technology Architect at Infosys, with more than three decades of experience in IT advisory, enterprise technology transformation, and managed operations. Vivin has led large-scale modernization programs for global enterprises, helping organizations strengthen operational resilience, accelerate transformation, and translate technology investment into business value.



External Document © 2026 Infosys Limited

For more information, contact askus@infosys.com

Navigate your next

© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.