Futurum®

NETSCOUT

Data Fitness for AI-Driven Operations:

Why Observability Alone is Insufficient When AI Operates at Machine Speed, Pushing Beyond Cloud Native's Heritage

APRIL 2026

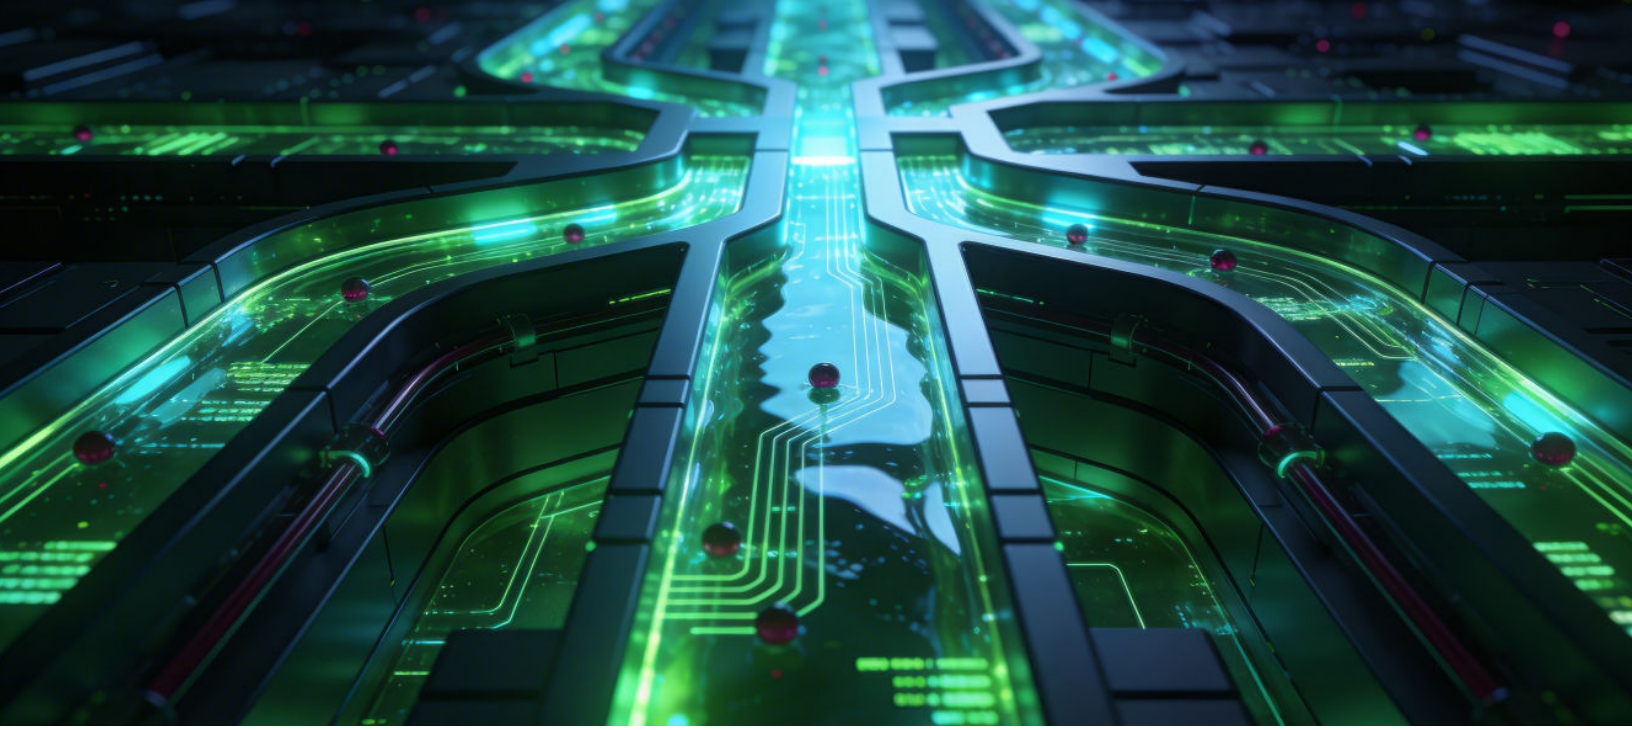

The Machine-Speed Challenge

Telemetry designed for human investigation will not support artificial intelligence (AI) acting autonomously at machine speed, as AI moves from analytical support toward autonomous execution. AI systems already detect anomalies and correlate signals; the operational future is executing remediation actions without human approval to compress decision cycles from hours to milliseconds. This shift creates a data fitness problem that organizations must address before expanding AI autonomy.

The Challenge: Today's observability telemetry designed for human investigation operates at human speed. Engineers navigate multiple tools, correlate data across fragmented sources, and reconstruct what happened from incomplete signals. Investigation unfolds over minutes to hours as operators compensate for sampling gaps, interpret aggregated metrics, and make inferences when information is incomplete or contradictory across tools.

The Need: Agents operate differently – they act on signals provided without the judgment layer humans apply. When telemetry loses fidelity through sampling or abstraction, AI cannot reliably interpolate what is missing – resulting in errors or worse, incorrect actions. Data confidence requirements must increase as decision latency decreases.

The business consequences are measurable, not theoretical: faulty responses trigger unnecessary interventions; incorrect root cause analysis extends mean time to resolution; and automated actions taken on incomplete data create revenue-impacting outages, compliance violations, and reputation damage. Forward-looking organizations must work now to address these issues.

Understanding what AI systems are doing and what they did requires data with the fidelity to reconstruct decisions, actions, and outcomes. Futurum's 2026 Software Lifecycle Engineering Decision-Maker Survey shows 37.4% of organizations now prioritize AI observability in platform selection, while 30.9% prioritize AI agent observability[1] (see Figure 1). Traditional categories remain foundational (cloud monitoring, security monitoring, and APM) but organizations recognize these alone are insufficient as AI autonomy expands. The gap between current observability capabilities and AI requirements is becoming visible.

1 Futurum Research (2026). Software Lifecycle Engineering Decision-Maker Survey

AI observability addresses two interrelated questions: monitoring how AI systems behave after they act, and ensuring the data AI consumes supports confident action before it acts. Data fitness is the input quality question that precedes the observability question.

Figure 1. Observability of AI Systems and Applications Emerges as a Core Platform Selection Criterion

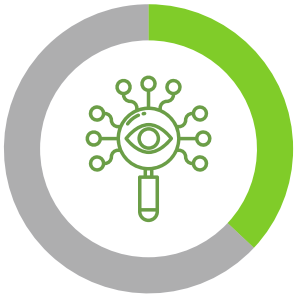

37%

of organizations prioritize AI observability in platform selections.

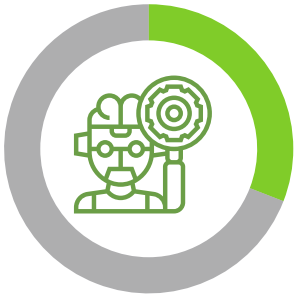

30.9%

of organizations prioritize AI agent observability

Futurum's 2025 Cybersecurity Decision-Maker survey shows that network traffic visibility, including encrypted traffic monitoring, is among the top three security priorities for the next 18 months

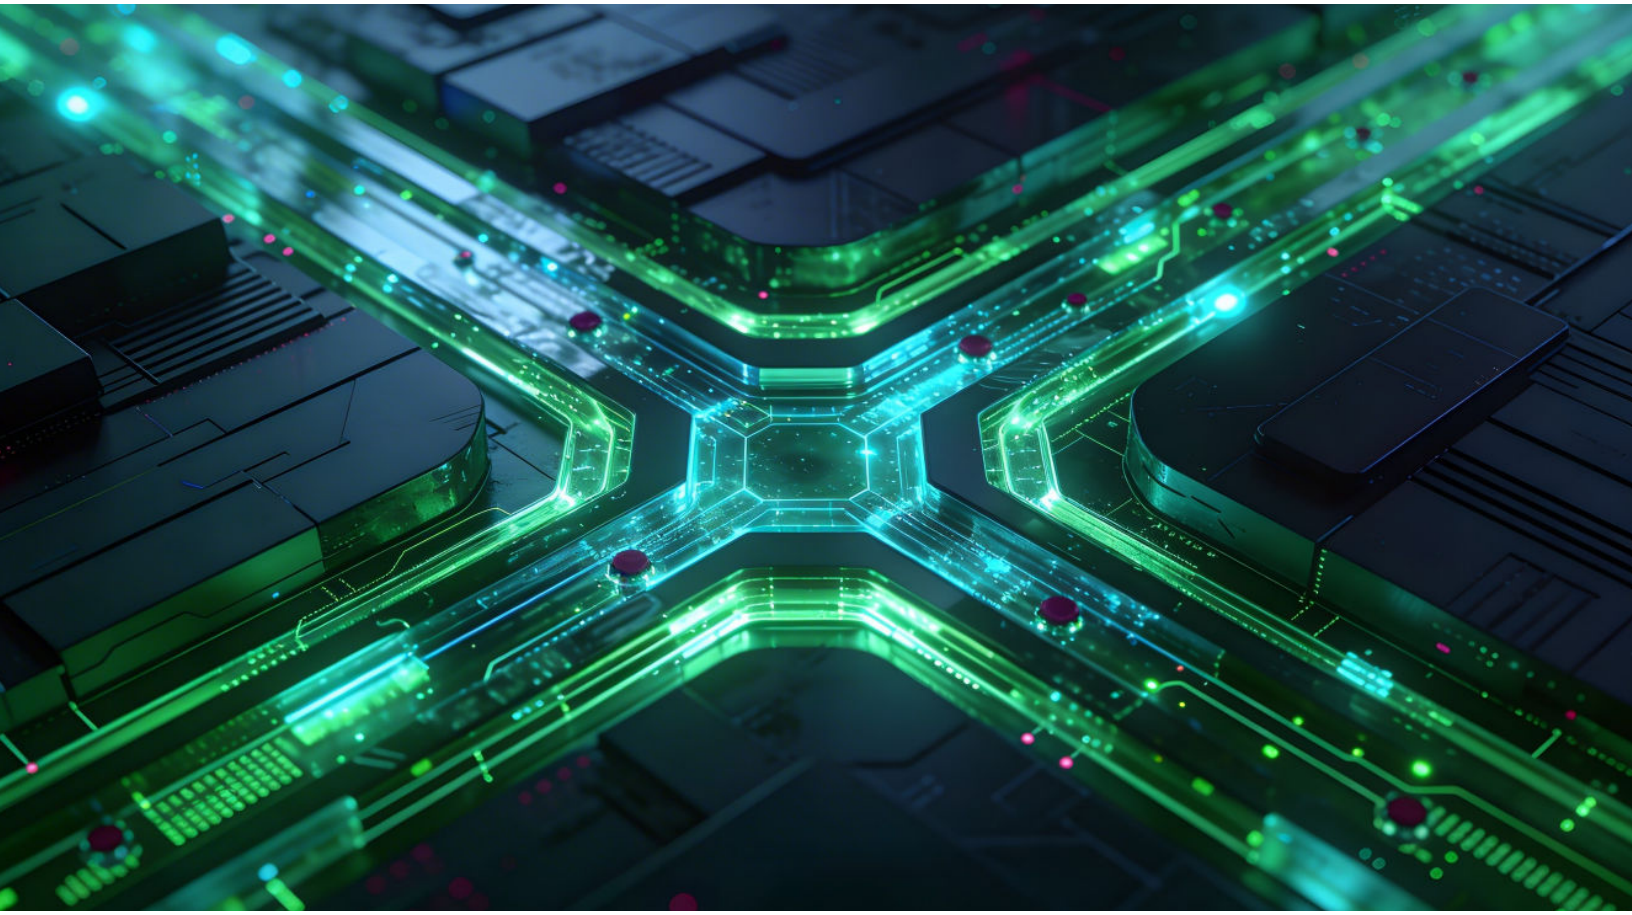

What Data Fitness Means

Data fitness describes input quality for AI-assisted and autonomous workflows: the confidence that data consumed by AI supports reliable action, not just retrospective analysis of what happened after AI acted. It is contextual, not absolute. Fitness depends on autonomy level, blast radius, operational maturity, and execution velocity. It answers the question: "Is this data sufficient for AI to act upon at this level of autonomy with accountability?"

Measurement Comparison for Data Quality, Observability Maturity, and Data Fitness

Six core attributes define fit-for-purpose operational data.

- **Fidelity:** Observed interactions provide higher confidence than instrumented abstractions. Actual system behavior carries more weight than telemetry filtered through application logging decisions. When AI acts on abstracted data, it inherits any embedded assumptions that may not hold under novel conditions.
- **Timeliness:** AI needs real-time data without gaps it cannot interpolate. Delayed or sampled telemetry forces AI to act on stale information or guess what happened between observations. Gaps tolerable for human investigation become blind spots when autonomous systems execute at machine speed.
- **Context:** Causality and sequence is essential; a snapshot of where things stand (state alone) is insufficient for AI to understand what led to current conditions and in what order. Without sequence, AI cannot distinguish cause from correlation to determine which action will resolve the problem – rather than compound it.
- **Provenance:** Verifiable chains establish accountability and auditability for tracing what data the agent observed, its decision, and why. As agent execution authority increases, provenance shifts from operational convenience to a regulatory and liability requirement.

What works for AI-assisted operations with humans in the loop may be insufficient when AI agents execute autonomously without approval gates.

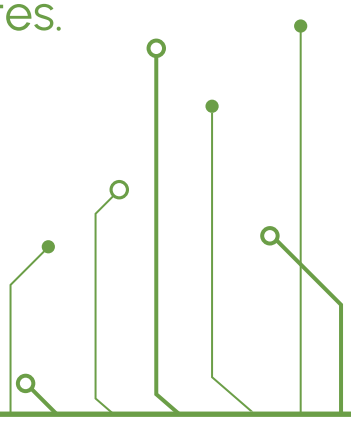

- **Consistency:** Coherence across observability, security, and network views is essential to prevent conflicting AI conclusions and eroding trust. Unified data sources prevent AI from taking actions that security flagged as violations or operations marked as unsafe.

- **Coverage**: Comprehensive visibility must not depend on anticipating behavior. AI must observe novel patterns and edge cases developers never instrumented; instrumentation bias creates blind spots at unexpected failure nodes – precisely where visibility matters most.

Traditional observability models embedded intentional tradeoffs; sampling controls storage costs; aggregation enables scale; and instrumentation focuses on anticipated behavior. These design choices were necessary when humans interpreted telemetry and compensated for what abstractions removed.

However, **as organizations introduce AI-driven workflows, these traditional observability tradeoffs create risk** by narrowing the acceptable margin for error:

Sampling creates visibility gaps AI cannot fill at machine speed without human judgement.

Aggregated metrics obscure the causality AI needs to attribute cause and effect.

Instrumentation bias means AI only sees what developers anticipated, missing novel patterns and edge cases that matter most.

The gap becomes acute as AI agents take operational roles. Agents make decisions, coordinate workflows, and execute actions across infrastructure spanning cloud providers, on-premises systems, SaaS platforms, and edge locations. Application instrumentation shows what each agent logs internally: decisions made, actions taken, outcomes reported. What it does not show is how agents actually interact with the systems they control and with each other.

The core challenge: Application instrumentation shows what agents log internally. Network-derived telemetry shows how agents actually interact with systems and each other.

Consider agent-to-agent communication, agent access to encrypted services, and agent interactions with platforms that teams do not instrument: application-centric observability struggles to capture these interaction layers where agent decisions will propagate and where failures in agent coordination will create cascading operational risk.

This interaction visibility gap represents the core data fitness challenge. Network-derived telemetry reduces this gap by observing actual system interactions at the wire level, independent of application instrumentation. Where application telemetry shows what developers exposed, network-derived approaches capture how systems actually communicate.

This complementary perspective becomes essential as AI agents operate across infrastructure boundaries where instrumentation fragments.

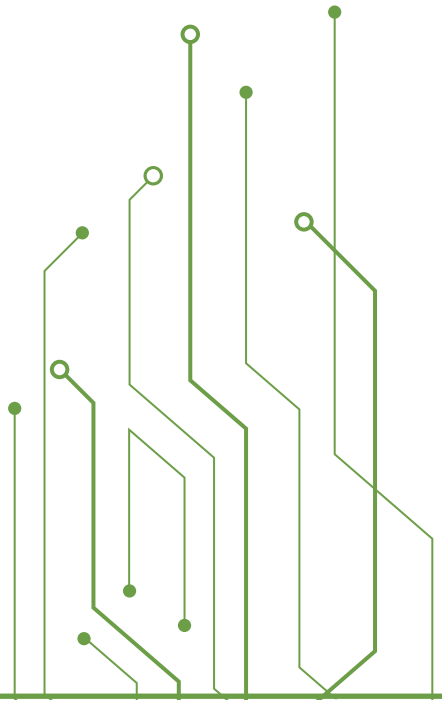

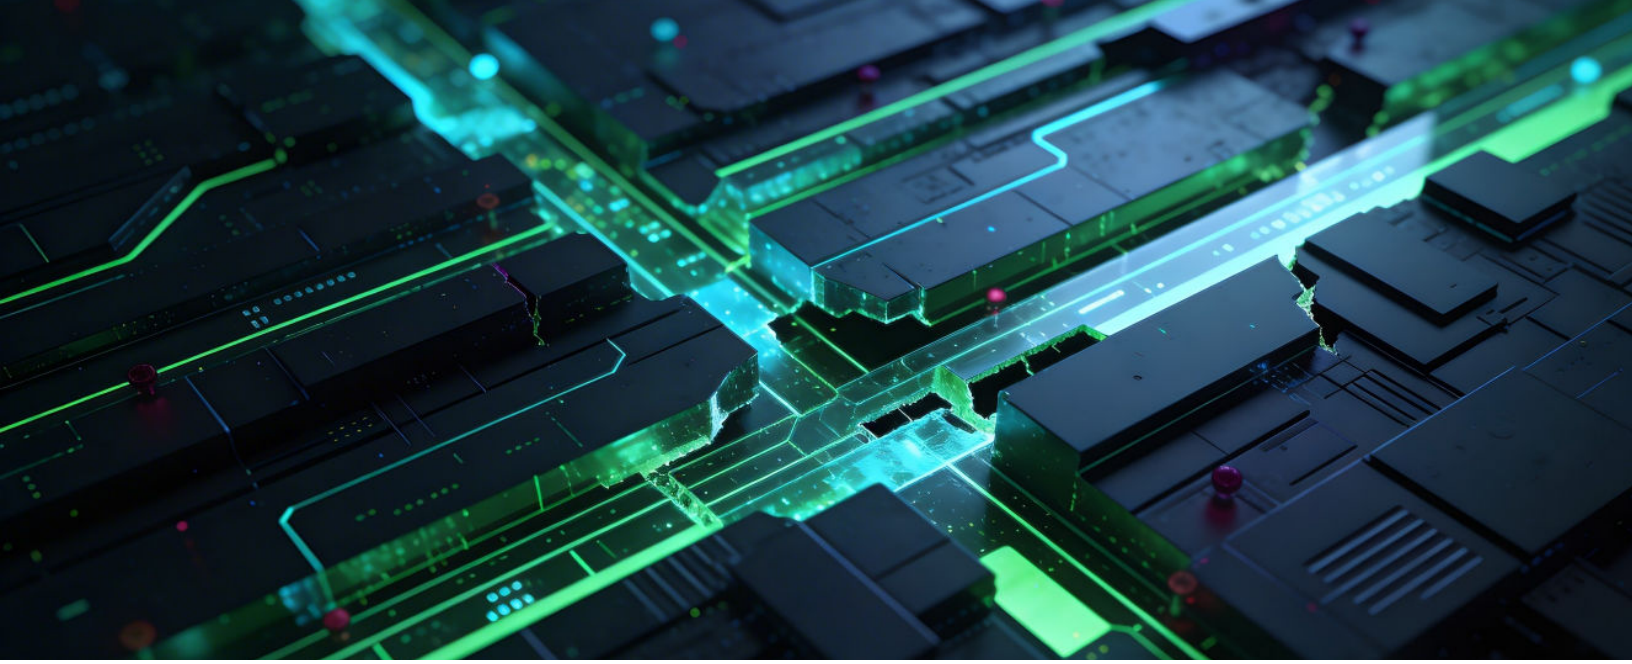

Strategic Decision Framework

Organizational AI readiness requires matching telemetry confidence to workflow autonomy levels: higher autonomy demands higher confidence. The framework must operate on graduated autonomy:

- **AI-assisted operations** keep humans in the loop, with AI systems recommending actions, assembling evidence, and adjusting thresholds.
- **Agentic operations** remove immediate human intervention as AI agents execute bounded actions or coordinate workflows autonomously. Telemetry becomes part of the control loop, not just analysis input. Data that is effective for assisted decision-making may require additional context when used to trigger automated action.

Data fitness becomes a gating consideration for where and how autonomous behaviors are introduced. As autonomy increases, operational data must support clear attribution of cause and effect, consistent interpretation across domains, confidence thresholds aligned to blast radius, and traceability establishing agent accountability.

How NETSCOUT Addresses the Interaction Visibility Gap

NETSCOUT generates network-derived telemetry through distilled metadata extraction. The technology observes actual interactions on the wire: which services communicated, when, in what sequence, with what protocols, and response characteristics.

This provides ground truth independent of whether applications are instrumented. SaaS platforms teams do not control service mesh encryption or legacy systems; network-derived telemetry sees communication patterns across boundaries where application instrumentation fragments.

NETSCOUT differentiates by extracting distilled metadata, rather than raw packet retention. The approach is purpose-built data feeds: curated, deduplicated, and optimized for cost-efficiency. This extracts interaction intelligence at scale without proportional cost growth. Coverage spans any cloud, any service, any application, any network, any vendor.

How NETSCOUT Addresses Data Fitness Requirements.

Fidelity:
Captures observed interactions, not abstractions.

Timeliness and Continuity:
Provides complete view without sampling gaps.

Context:
Demonstrates sequence and causality across service boundaries.

Provenance:
Captures verifiable network behavior.

Consistency
Establishes a canonical data source applicable across observability, security, and operations teams.

Coverage
Ensures comprehensive visibility regardless of instrumentation.

For AI-driven operations, AI observes behavior it will act upon (actual interactions, not inferred abstractions), decreasing ambiguity in the telemetry that feeds autonomous decisions.

The positioning is complementary: Application telemetry provides code-level context, business logic, and user journey visibility while network-derived telemetry provides interaction context that application instrumentation cannot see.

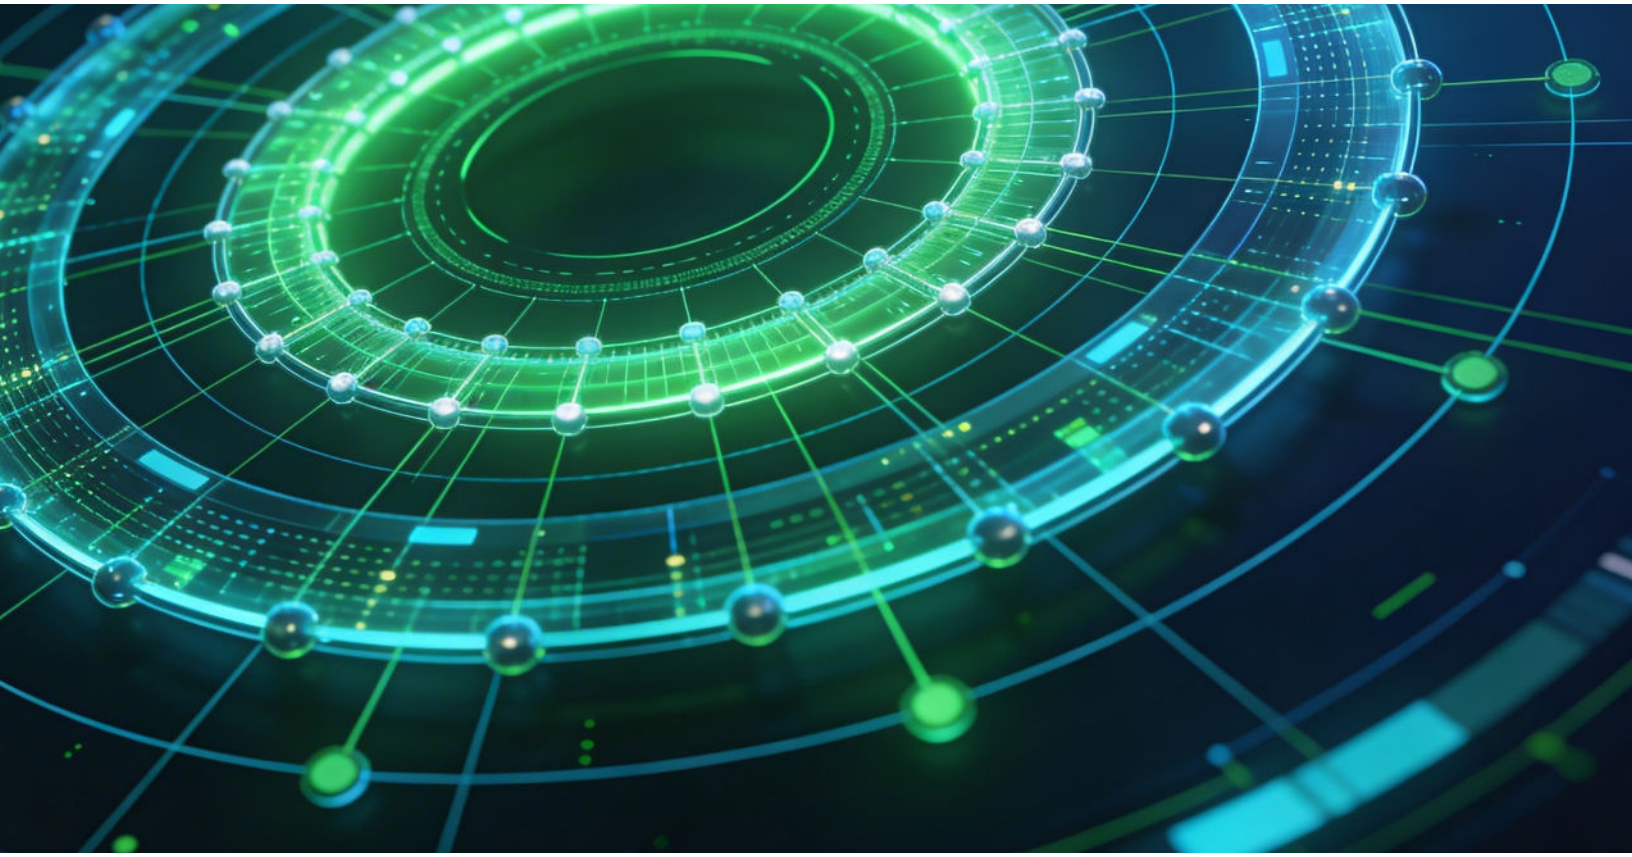

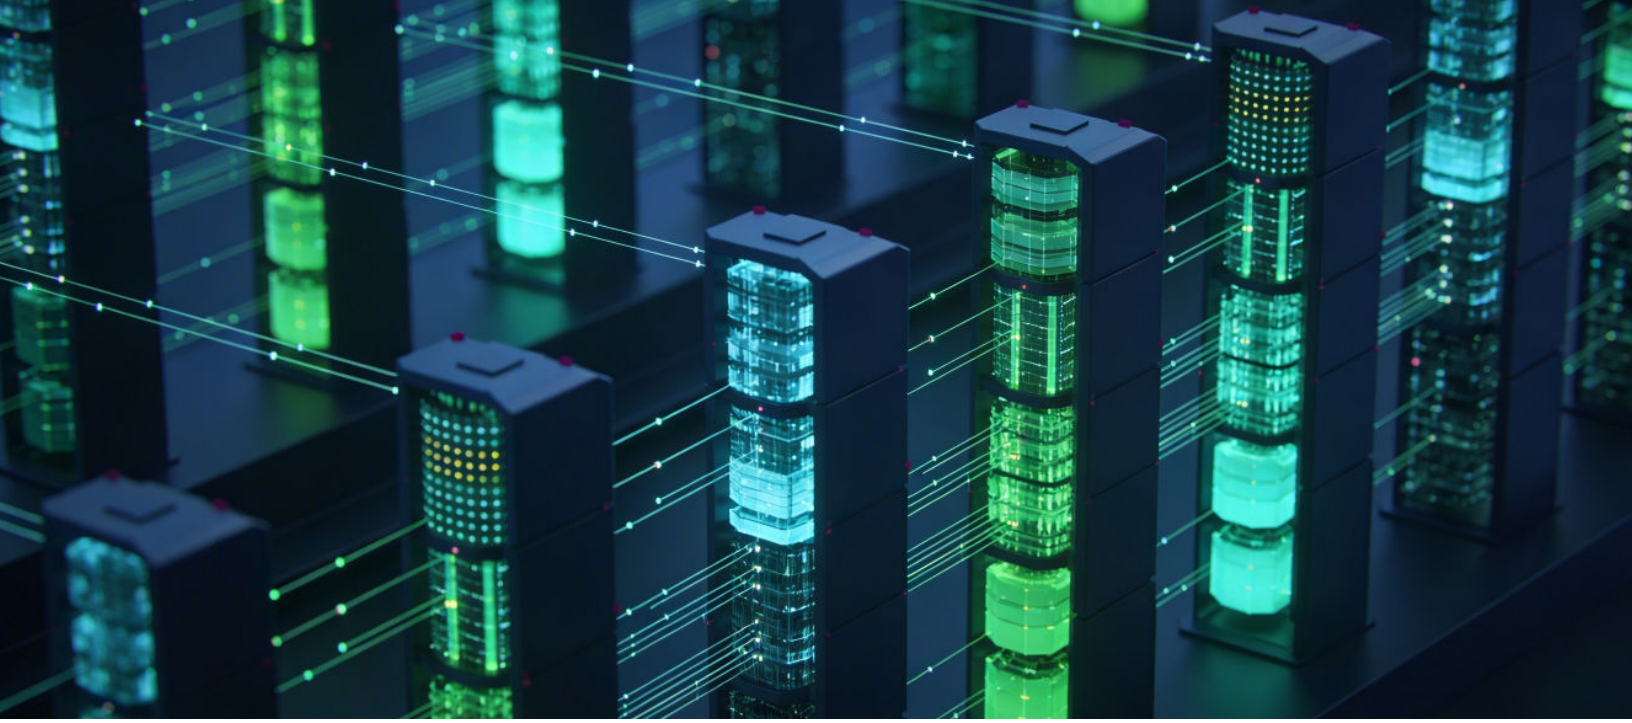

Conclusion

AI adoption velocity will be constrained by telemetry confidence. Data foundations built for human interpretation at human speed require reassessment as AI moves toward autonomous execution at machine speed.

Organizational readiness requires data fitness as a strategic component. Organizations that align telemetry strategy with automation scope and autonomy levels will apply AI more responsibly and effectively than those that assume existing observability investments automatically support machine-speed operations.

The convergence of observability, security, and AIOps demands unified data approaches. Data fitness is the defining factor in operational maturity versus operational risk.

The strategic question for technology leaders: **Is your data fit for the decisions AI will make?**

Important Information About This Report

AUTHORS

Mitch Ashley
Vice President & Practice Lead Software
Lifecycle Engineering | The Futurum Group

Fernando Montenegro
Vice President & Practice Lead,
Cybersecurity & Resilience | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

NETSCOUT

ABOUT NETSCOUT

NETSCOUT Systems is a provider of service assurance, cybersecurity, and business analytics solutions designed to help enterprises and service providers maintain the availability, performance, and security of their digital environments. The company is best known for its deep packet inspection technology and its nGenius® platform, which delivers real-time visibility into network and application performance across complex, distributed infrastructures. NETSCOUT's capabilities extend into DDoS protection and threat intelligence through its Arbor® portfolio, enabling organizations to detect and mitigate cyber threats at scale. With a focus on high-fidelity data and operational insight, NETSCOUT supports organizations in reducing downtime, improving user experience, and strengthening overall digital resilience.

Futurum®

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum®

CONTACT INFORMATION: **The Futurum Group LLC | futurumgroup.com | (833) 722-5337**