HyperFRAME RESEARCH

RESEARCH BRIEF

Closing the Observability Gap for Today's Distributed Applications

Why Internet Performance Monitoring Is Essential for Modern Applications

APRIL 2026

Author:

Stephanie Walter
Practice Leader, AI Stack

Steven Dickens
CEO & Principal Analyst

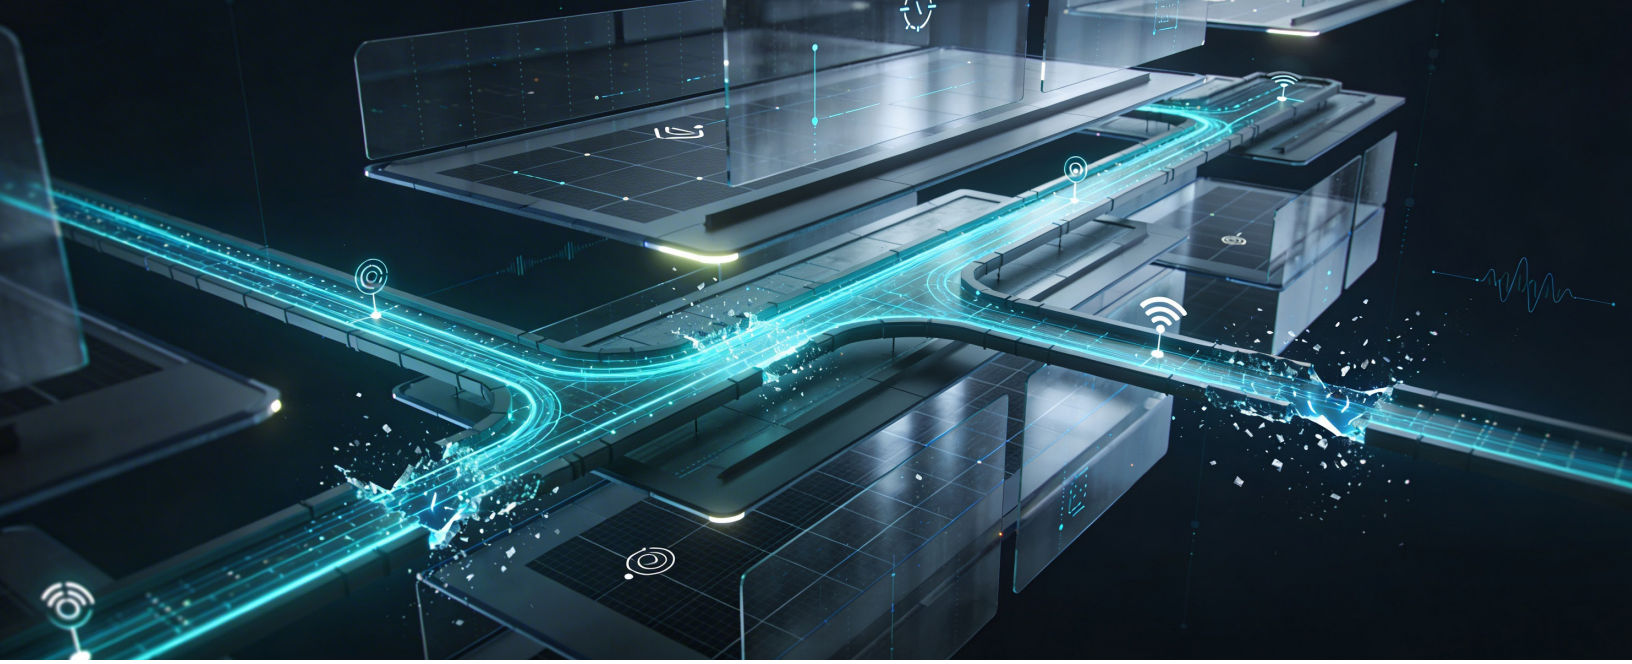

Executive Summary

Enterprise outages increasingly occur outside the boundaries of internal infrastructure. The disconnect between internal visibility and real user experience represents a growing observability gap in modern distributed environments. Modern enterprise applications now exist as a distributed fabric of cloud platforms, SaaS services, third-party APIs, and internet infrastructure. While traditional observability focuses on internal health, it often ignores the internet stack, including the DNS, CDNs, ISPs, APIs, WANs, cloud services, and routing infrastructure that delivers these services. These dependencies are increasingly the primary source of performance failures. This risk is amplified by AI systems and agentic workflows that require constant, low-latency connectivity to services. Without internet-native telemetry, AI-driven operations are incomplete and unable to correlate code performance with the broader set of services, APIs, and networks applications rely on.

Internet Performance Monitoring (IPM) addresses this by measuring how internet infrastructure affects application delivery. By leveraging a globally distributed observability network, organizations gain visibility into real-world application or service delivery where users are, while also understanding routing and network performance globally. Integrating IPM with hybrid observability platforms allows for full-path visibility that synthetic testing alone cannot deliver without integration into infrastructure observability workflows. This architectural visibility cannot be replicated through endpoint agents or synthetic testing alone, as these methods lack continuous, real-world network context.

The Distributed Reality of Modern Applications

The shift from monolithic architectures to distributed services has fundamentally changed the nature of IT operations. Today, a single user transaction might traverse multiple cloud authentication and security systems, edge infrastructure locations, and API services before reaching its destination. Our analysis suggests that the number of these dependencies is set to expand as AI adoption moves from experimentation to production. According to the HyperFRAME Research Lens, while only 30% of organizations have AI deployments at scale today, 66% expect mass AI deployments within 12 to 24 months. This rapid expansion creates a massive web of external system dependencies that traditional monitoring was never architected to handle.

Enterprise environments are already deeply fragmented. HyperFRAME Research Lens data shows that 37% of organizations operate hybrid data architectures that combine legacy systems with cloud platforms. This creates a complex reality of data movement where service dependencies cross multiple jurisdictional and network boundaries. Every jump between a data center, cloud provider, office, remote location, and edge introduces a point of failure that exists outside the enterprise network perimeter.

Increasingly, core business functions are delivered through third-party services rather than internal systems. Payments, authentication, communications, analytics, and AI capabilities are frequently executed through SaaS platforms and external

APIs, creating dependencies that extend far beyond enterprise-controlled infrastructure.

The digital delivery chain now includes invisible links like DNS resolution, CDN routing, ISP backbone networks, APIs, and identity management systems. If any of these external services falter, the application appears down to the user, even if internal server metrics look healthy. Increasingly, performance failures originate outside enterprise-controlled infrastructure, making external dependencies the dominant failure domain for modern applications.

The pattern is well established. Over the past decade, large-scale outages tied to DNS failures, routing errors, and widely distributed software updates have repeatedly disrupted global services despite healthy internal infrastructure metrics. In these scenarios, internal dashboards often reported normal system performance while users experienced widespread service disruption. Organizations with external visibility into internet routing, DNS infrastructure, and service dependencies can detect the underlying pattern earlier than those relying solely on internal telemetry.

Furthermore, modern AI architectures are leaning heavily into Retrieval-Augmented Generation (RAG). HyperFRAME reports that 78% of organizations plan to implement RAG within the next 12 months. This introduces even more reliance on external data services and API-driven context pipelines. When IT teams lack visibility into these external segments, they create significant operational blind spots. Traditional monitoring tools were originally designed for environments where most systems operated within enterprise-controlled infrastructure. As application delivery moved into distributed, internet-dependent ecosystems, these tools were not architected to observe the full external delivery path. We are moving into a world where the majority of the code you run and the networks you use are not yours.

Why AI Systems Increase Operational Fragility

AI systems are often more operationally sensitive than traditional applications because they are architected to be highly synchronous and data-intensive. An AI inference pipeline often requires real-time access to model endpoints, vector databases, orchestration frameworks, interface layers such as MCP, and security enforcement services. If there is a spike in latency in a DNS lookup, API response, or a hiccup in an internet routing path, the entire AI workflow can time out or provide a degraded experience.

Agentic workflows represent the next level of this complexity. These agents do not just perform a single task. They chain together multiple API calls, retrieval steps, and cross-service orchestrations. Many of these workflows rely on Retrieval-Augmented Generation (RAG) pipelines, where sequential inference and external data retrieval must occur in tightly coordinated steps. Every link in that chain must succeed for the workflow to complete. A failure in a third-party API or a slow response from a remote vector store does not just delay the process. It often breaks the logic of the agent entirely. In many cases, security validation layers such as identity checks and policy enforcement also sit inline with these workflows, introducing additional points where latency or failure can cascade across the pipeline.

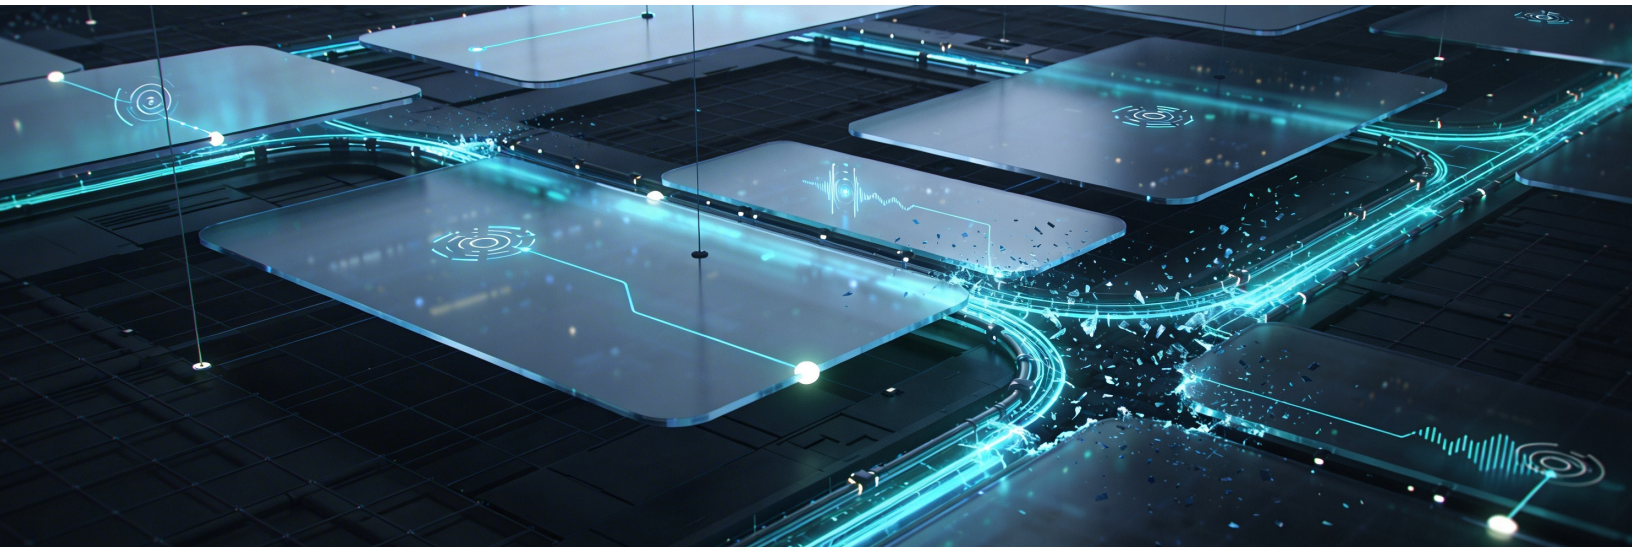

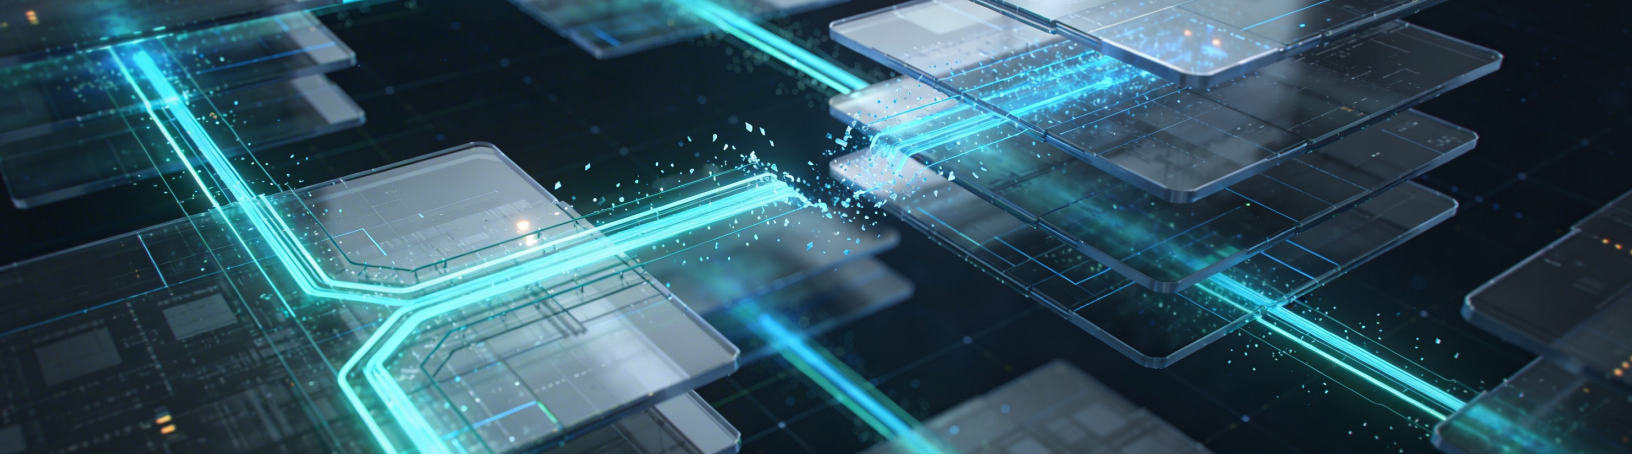

This creates a strategic tension for IT leadership. On one hand, the business demands the speed and capability of AI agents. On the other hand, the operational complexity of managing those agents across the public internet is staggering. Standard infrastructure monitoring tools focus on code, internal logs, CPU, and memory. They are not designed to see the Border Gateway Protocol (BGP) route leak that is slowing down your API calls to an LLM provider. To maintain any level of reliability, observability strategies must expand to include internet-level performance visibility. Without it, the brain of AI is disconnected from the body of enterprise data.

Think of it this way: AI is an amplifier. It amplifies whatever data pipeline it is connected to, good or bad. An AI that trains on telemetry with a seventy percent blind spot does not just miss things. It optimizes confidently for what it can see. The backend infrastructure may look flawless, but users can still suffer. And the AI will show everything is fine. This is no longer theoretical. As enterprise AI workflows expand across distributed services, failures in shared infrastructure and external dependencies are increasingly cascading across interconnected systems. The enthusiasm around AI adoption is accelerating, but optimism without visibility remains a structural risk. Fix the data first. Then unleash the AI.

What Internet Performance Monitoring Actually Measures

Internet Performance Monitoring aims to deliver a comprehensive view of the health, performance, and resilience of every service or system in the network path between the user and the application across the internet infrastructure. It goes beyond checking if a site is up and instead tracks the health of the entire internet stack. This includes DNS resolution times, CDN routing efficiency, ISP connectivity for every office or retail location, and the availability and latency of internal and external APIs, internal networks, WAN, and connectivity between locations and data centers that modern applications depend on.

Effective IPM combines several measurement techniques to provide a full picture. It uses synthetic testing from distributed global agents to simulate user journeys and Real User Monitoring (RUM) to capture actual performance in real time. It also tracks BGP route changes, hijacking attempts, and backbone routing paths.

Unlike internal tools, IPM measures performance from global vantage points that simulate real user experiences. These vantage points span thousands of autonomous systems and ISP environments, enabling organizations to observe how services behave across real-world internet routing paths. This allows a team to determine if a performance problem is local to a specific region, a specific ISP, or if it is a global infrastructure failure. This level of internet-native telemetry is impossible to gather from inside your own data center or even from within a cloud VPC. It provides the ground truth. When a problem occurs, the first question is always whether it is an internal error or an external dependency. IPM enables teams to diagnose outages faster and with greater accuracy, determining whether failures originate within internal infrastructure or external dependencies.

Ultimately, the purpose of internet visibility is not simply to optimize infrastructure performance but to understand how users experience business-critical services. IT organizations are increasingly measured by the reliability and responsiveness delivered to employees, customers, partners, and automated systems. This shift is driving greater emphasis on user-experience-centric metrics, including experience-level objectives (XLOs) and AI-assisted experience scoring that correlate technical performance with real business impact. Without visibility into user experience across the internet path, organizations cannot accurately determine whether systems are meeting operational expectations or delivering acceptable service outcomes.

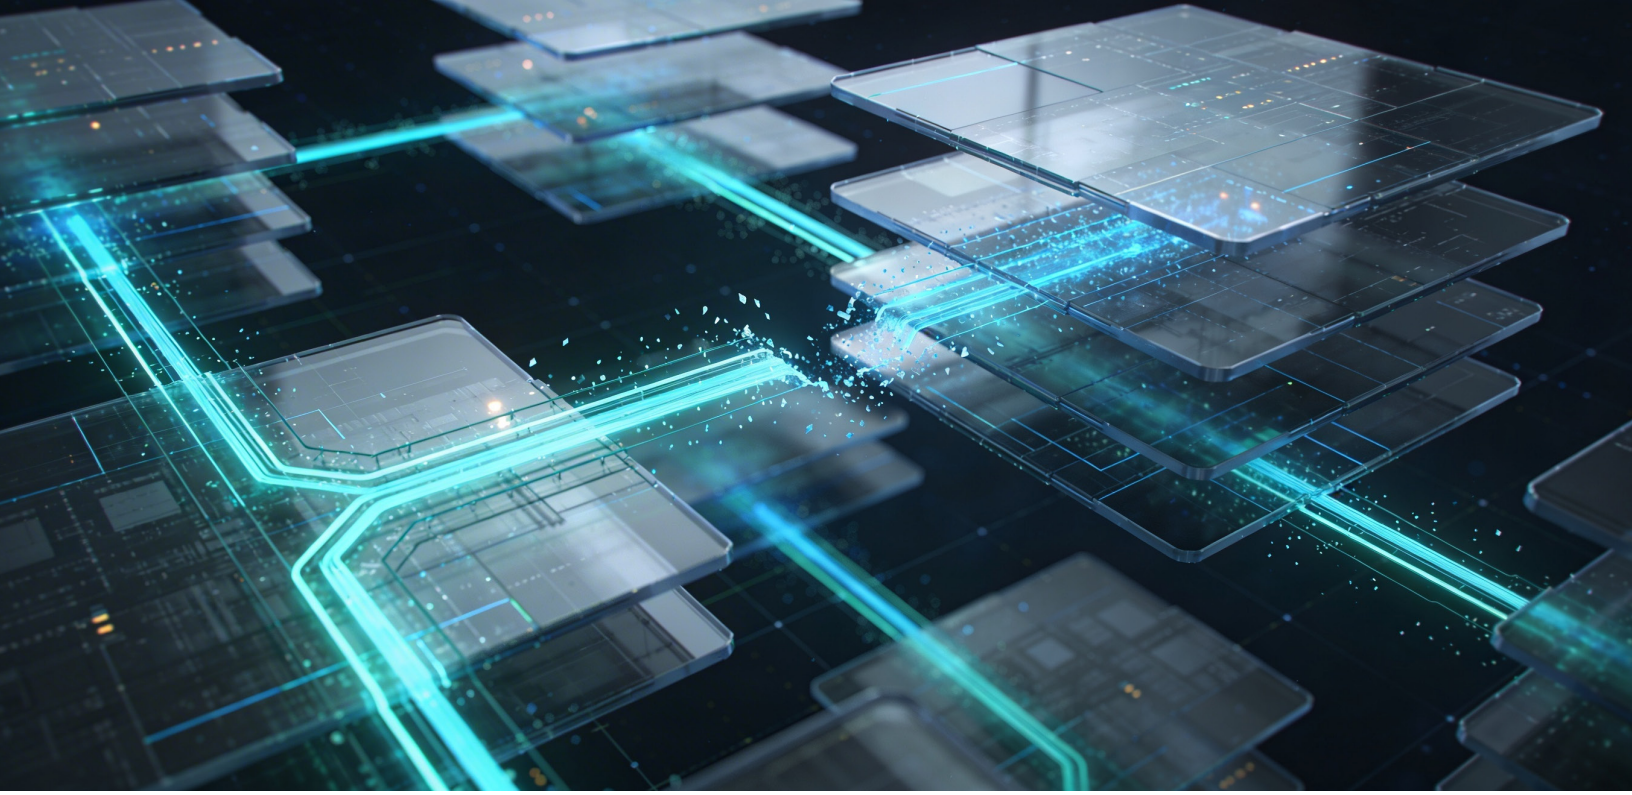

Recommendations for an Internet Performance Monitoring Platform

Selecting the right Internet Performance Monitoring (IPM) platform is a strategic infrastructure decision, not simply a tooling choice. As organizations expand their reliance on distributed applications, external APIs, and SaaS platforms, the ability to observe internet dependencies becomes essential to maintaining reliability and performance. When evaluating an IPM platform, leadership must look past the marketing and focus on the architecture of the monitoring network. We recommend prioritizing the following criteria:

- **Global vantage point coverage:** Ensure the provider operates a purpose-built network distributed across real consumer ISPs and backbone providers rather than relying on limited third-party infrastructure. Organizations should align monitoring locations with their office sites, customer geographies, and critical service regions to ensure coverage reflects real user experience. True internet visibility depends on measuring performance from the same networks users rely on.
- **Internet stack visibility:** The platform must provide deep metrics across the full internet delivery chain, including DNS resolution, CDN routing, ISP connectivity, API performance, MQTT, MCP, TLS, SSL, SIP, and backbone routing behavior. Without visibility into these components, teams cannot fully diagnose service degradation.
- **Integration with observability platforms:** Data should not sit in a silo. IPM telemetry must integrate directly with infrastructure monitoring and operations workflows to ensure unified investigation and response processes.
- **Root-cause intelligence:** The system should be architected to distinguish between internal application failures and external network events using dependency mapping and AI-assisted root-cause intelligence capabilities.
- **Automation readiness:** Monitoring insights should be accessible via APIs and workflows that enable automated investigation and remediation, reducing the time between detection and resolution.
- **Business risk alignment:** The platform should support visibility that maps technical performance to revenue, customer experience, and regulatory exposure.

Organizations that prioritize these capabilities reduce operational blind spots, improve diagnostic accuracy, and build digital resilience in environments increasingly dependent on external internet infrastructure.

When Organizations Should Prioritize Internet Performance Monitoring

Not every application requires deep internet visibility, but for most modern enterprises, the critical list is growing. We recommend prioritizing IPM if digital revenue or customer experience is tied directly to web performance. Organizations whose applications rely heavily on third-party APIs or external service integrations should also treat internet visibility as a critical operational requirement. If a 100ms delay in API response time costs thousands of dollars, the internet path needs to be seen. In many industries, degraded digital performance also translates directly into brand damage and loss of customer trust.

Organizations with a geographically distributed workforce also face unique risks. This risk increases significantly in multi-cloud environments where applications span multiple providers and regions beyond traditional network boundaries. If employees rely on SaaS platforms for daily productivity, their office is the internet. When they experience lag, your internal help desk is often powerless without IPM data. Organizations operating highly distributed infrastructure and relying on external service dependencies face increasing operational risk without visibility into the internet stack. Furthermore, as AI-driven services move into regulated industries, the cost of a reliability failure increases. Modern AI workflows depend on distributed model endpoints, retrieval pipelines, and external data services, all of which introduce additional points of failure across the public internet. In these cases, IPM is not just about performance. It is about compliance and risk management. If you cannot prove where a failure happened, you own the liability for it. Many traditional tools rely on internal telemetry, which limits their ability to diagnose these external issues.

Bridging Internet and Infrastructure Observability

The real value of observability is found in the correlation of disparate data sets. Hybrid observability platforms excel at monitoring the servers and the code. IPM excels at monitoring the delivery path. Bringing them together creates a unified operational context.

The integration between Catchpoint and LogicMonitor is a prime example of this unified architecture for full-path observability. It is not just about a simple data exchange. It is an architectural alignment. LogicMonitor provides the deep visibility into the hybrid infrastructure on-premises and in the cloud. Catchpoint adds the external internet telemetry.

The architecture maps to the problem. LogicMonitor covers the thirty percent organizations own and control: applications, servers, cloud infrastructure, containers, and internal services. Catchpoint covers the seventy percent organizations do not own but depend on: DNS resolution, CDN edge behavior, ISP routing, last-mile connectivity, BGP path changes, and the backbone networks that traffic traverses before it reaches any user. Together, they give Edwin AI a complete dataset. With that completeness, an AI system can make more trustworthy decisions.

When you layer an AI-driven system like LogicMonitor's Edwin AI on top of this combined dataset, the results can

be transformative. The AI can see that a spike in application latency in LogicMonitor correlates exactly with a BGP routing change detected by Catchpoint. This allows for near-instant causality determination. Instead of a war room with twenty people guessing, what the industry calls Mean Time to Innocence, or the costly process of ruling out every internal system before finding an external cause, there's a single source of truth that points to the specific ISP or service provider at fault. This model aims to deliver a closed-loop system where the AI can eventually initiate automated remediation. Full-path observability is the only way to manage the modern digital delivery chain effectively.

Looking Forward

Digital services are increasingly delivered across distributed cloud infrastructure and global internet networks. As AI-enabled applications and agents become more autonomous, they will rely on a web of interconnected services and APIs that no single company owns. This makes the internet the most critical, yet least understood, part of the modern tech stack.

Observability platforms must evolve. Traditional infrastructure monitoring is table stakes. The next generation of tools will be defined by their ability to provide full-path visibility. Internet stack visibility is becoming a foundational requirement for IT operations. Organizations that continue to ignore the internet-native telemetry domain will find themselves perpetually reactive.

We believe that the convergence of infrastructure observability, internet intelligence, and agentic AI will define the next generation of operational tooling. Agentic AIOps requires high-fidelity data to make correct decisions. Without internet-native visibility, these systems are operating with a massive blind spot. The winners will be the organizations that adopt a persistent, global observability fabric to act as their external control plane. Organizations that adopt this integrated model will be better positioned to maintain resilience, diagnose outages quickly, and deliver reliable digital experiences. The goal is no longer just to see your systems. The goal is to see the entire world they live in. Observability vendors will increasingly converge toward unified, full-path architectures that combine infrastructure telemetry, internet intelligence, and agentic AI into a single operational system. Over time, organizations that treat internet visibility as optional will face escalating reliability risks, while those that operationalize full-path observability will define the next generation of resilient digital enterprises.

ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

CONTACT HYPERFRAME RESEARCH:

Steven Dickens
CEO & Principal Analyst | HyperFRAME Research

Email Address:

steven.dickens@hyperframeresearch.com

Telephone Number:

+1 845 505 1678

X: @StevenDickens3

LinkedIn: Steven Dickens

BlueSky: Steven Dickens

CONTRIBUTORS

Stephanie Walter
Practice Leader, AI Stack

INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research 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 "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

LICENSING

DISCLOSURES

HyperFRAME Research 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.


HyperFRAME
RESEARCH