NOVEMBER 2024

LogicMonitor Brings Foundation for High Impact AI to Hybrid Observability

Jim Frey, Principal Analyst, and Torsten Volk, Principal Analyst

Hybrid Observability as the New Minimal Requirement for Operations

As modern applications are often no longer tied to their underlying infrastructure, consist of numerous loosely coupled microservices, and can rapidly move and scale, corporate IT needs urgent help to understand and prioritize the potential impact of the yellow and red lights on their monitoring dashboards. This is where observability must come in, in place of traditional monitoring, to save the day. While monitoring is a mostly static discipline based on the collection of predefined metrics and logs—often separate from applications, infrastructure, and services to detect anomalies or failures—observability provides deeper insights into the key factors influencing application performance and resilience. In other words, monitoring keeps an eye on static, often disconnected metrics, while observability looks for unknown issues based on the holistic analysis of metrics, traces, and logs across application stacks. As a result, observability can alert organizations about problems across the continuously changing environment and context.

Today, a custom enterprise application might be running on a Dell/VMware data center stack, but tomorrow it could be distributed across multiple cloud platforms, leveraging containerized microservices on AWS, serverless functions on Azure, and machine learning services from Google Cloud. Instead of looking at static metrics attached to infrastructure components, observability must be aligned to the overall applications, independent of architecture, infrastructure, and location. This puts corporate IT in the driver's seat, as it can now prioritize proactive maintenance based on application impact rather than reactive firefighting tasks. Going one important step further, observability can map the impact of any infrastructure change back to the applications that it supports and all the way through to business KPIs. This enables IT operators to prioritize working on tasks based on business value, such as optimizing the transaction speed of an e-commerce shopping cart, over tasks that might not have the same degree of direct business impact (e.g., improving user experience on the corporate website).

As operational best practices evolve to deal with the steadily increasing complexity and scale of IT infrastructure and the applications they support, enterprises must move from traditional monitoring to proper observability approaches. And in today's hybrid, multi-cloud, and increasingly container-based environments, observability strategies must also be hybrid in nature to cover modern application infrastructures as well as critical legacy systems upon which many organizations still rely. To do anything less than to strive to be as comprehensive as possible from a cross-domain perspective puts operational integrity, user experience, and business success at unacceptable risk.

Making AI Work for Hybrid Observability

Various forms of AI, e.g., including generative AI, anomaly detection, predictive analytics, and reinforcement learning, represent exciting potential for improving the speed and effectiveness of operational management solutions. But AI does not replace the basic requirements for observability platforms; instead, it is meant to build upon those requirements. TechTarget's Enterprise Strategy Group believes that four key pillars must be addressed to set the proper foundation for effective application of AI: hybrid observability, intelligent data collection, auto-discovery, and dependency mapping (see Figure 1). These are all required to provide AI with the data needed to generate actionable insights and alerts, automated root cause analysis (RCA), accurate predictions, and contextualized responses to questions coming from IT users, developers, security engineers, or the line of business.

Figure 1. The Four Pillars of Comprehensive Observability

Source: Enterprise Strategy Group, a division of TechTarget, Inc.
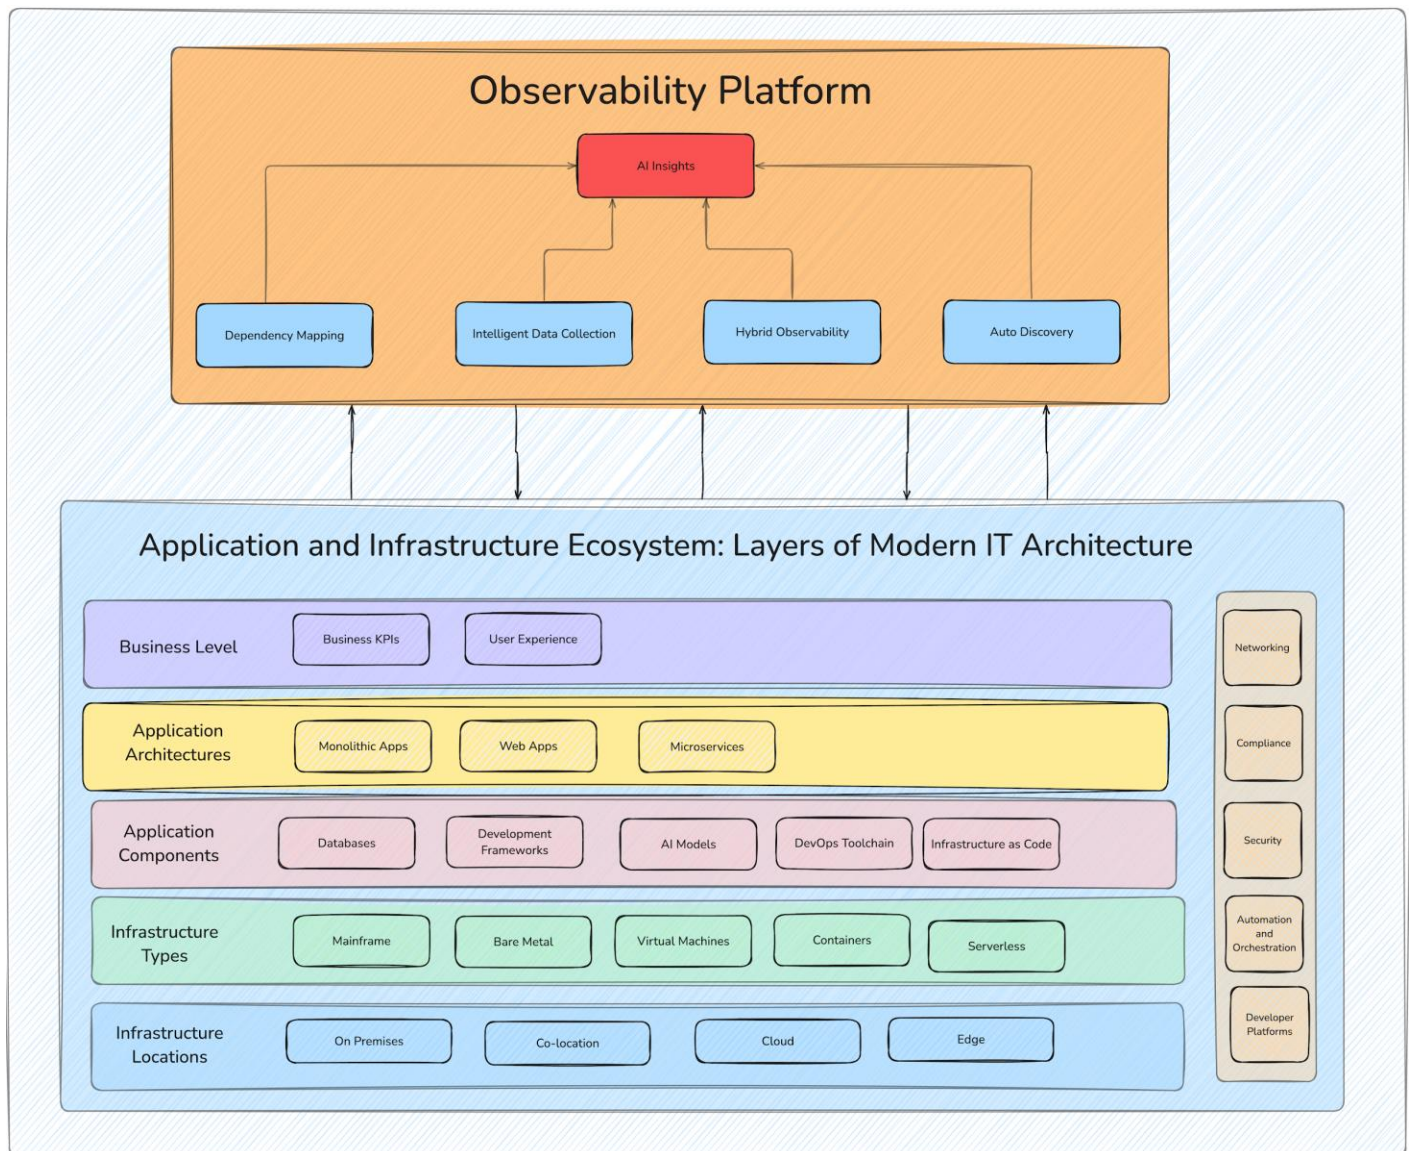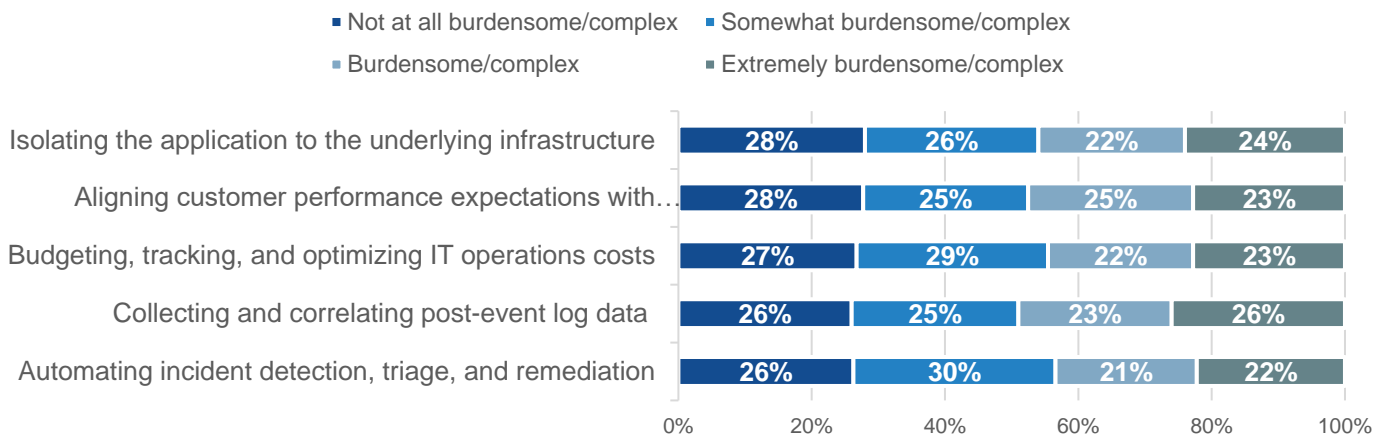
Enterprise Strategy Group
by TechTarget

The four pillars of observability create the foundation for optimal AI effectiveness in the following ways:

- **Hybrid observability.** By normalizing and unifying telemetry data across all infrastructure layers, AI gains a complete view of the system via elimination of data silos and blind spots.

- **Intelligent data collection.** Deep and consistent telemetry collection ensures that AI has high-quality, detailed, application-aware data for learning and inferencing.

- **Auto-discovery.** Automatically detecting and monitoring new or changed components ensures that AI is always analyzing the most up-to-date environment across both east-west and north-south dimensions.

- **Dependency mapping.** By continuously mapping relationships and dependencies across components, AI can better "understand" the blast radius and cascading effects of incidents, including user experience and service quality impacts.

If these needs are met, the use of AI becomes a powerful means of leveraged acceleration to accurately correlate data and events, improving RCA convergence, reducing time to restoration, and presenting realistic paths to proactively protecting application performance and user experience. Done properly, AI can help teams address their most problematic and burdensome tasks, as identified by Enterprise Strategy Group research in Figure 2. [1]

Figure 2. The Five Most Burdensome IT Operations Tasks

For each of the following operations, please rate the level of resource burden and overall complexity the operation creates for your organization's IT operations team. (Percent of respondents, N=360)

- Not at all burdensome/complex
- Somewhat burdensome/complex
- Burdensome/complex
- Extremely burdensome/complex

	Not at all	Somewhat	Burdensome	Extremely
Isolating the application to the underlying infrastructure	28%	26%	22%	24%
Aligning customer performance expectations with…	28%	25%	25%	23%
Budgeting, tracking, and optimizing IT operations costs	27%	29%	22%	23%
Collecting and correlating post-event log data	26%	25%	23%	26%
Automating incident detection, triage, and remediation	26%	30%	21%	22%

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

The LogicMonitor Approach

LogicMonitor's LM Envision observability platform places a strong focus on delivering the four pillars of AI-driven hybrid observability for any combination of application architectures and infrastructures on-premises, in colocation facilities, in the public cloud, or at edge locations. At its core, this ability to instantly adjust to complex brownfield environments is a direct result of LogicMonitor's experience in building interfaces and connectors to integrate with everything from monolithic applications tied to mainframe hardware all the way to distributed microservices applications running on multiple Kubernetes clusters at different data centers, clouds, or edge locations. LogicMonitor's 3,000+ integrations are the true star of the show, as they ensure the consistent collection of relevant

[1] Source: Enterprise Strategy Group Research Report, *Generative AI in IT Operations: Fueling the Next Wave of Modernization*, September 2024.

telemetry data, the creation of service incidents, and, ultimately, the triggering of automated remediation and resolution workflows.

LogicMonitor introduced solutions that started with network and systems metrics and event collection many years ago. Along the way, it has expanded its platform to include log analytics (via acquisition of Unamoly), application health (via acquisition of Airbrake), and, most recently, AI analytics (via acquisition of Dexda). Importantly, LogicMonitor hasn't just strapped these components to its base platform via duct tape and string; rather, it has taken the time to thoughtfully and thoroughly integrate them into its observability platform, resulting in a seamless experience for operators.

Some of LogicMonitor's latest extensions include the Resource Explorer and Service Insights features, which enable easy organization of observability data related to individual services, applications, and locations. These capabilities help focus IT operations teams on business-critical applications and infrastructure, so they can better communicate with line-of-business teams and nontechnical leadership as to the state of the IT environment. For example, if a critical e-commerce app experiences an unusual decrease in transaction volume, Service Insight can instantly connect the issue to a database service running in a specific cloud region. This enables corporate IT to alert stakeholders immediately and offer database admins actionable insights to quickly resolve the issue, while updating nontechnical leadership on the potential revenue impact and expected timeframe to resolution.

While LogicMonitor has been applying AI technologies to hybrid observability for years via features such as metric anomaly detection, dynamic thresholding, and pattern recognition, the addition of the Dexda large language model AI engine, now known as Edwin AI, marked an important turning point. Launched in June 2024, Edwin AI has high potential for a truly transformative impact due to its use with an established observability platform, LM Envision, that provides the four pillars described above.

LogicMonitor has first applied Edwin AI to the task of correlating and reducing the often-overwhelming volume of events that comes along with true and proper hybrid observability. As part of correlation, Edwin AI generates insights and points to relevant relationships across data types and source domains, helping accelerate causal analysis and root isolation.

And the results have been impressive. Multiple customers are seeing event reduction rates of more than 80%, enabling operations teams to focus on what matters instead of digging through piles and piles of events to find the important and relevant indicators. MSPs using the LM Envision platform to manage customer environments have seen that Level 1 operators can deal with more issues, reducing escalations and enabling more senior staff to remain focused on value-add projects.

There are future plans, of course, to further expand the ways in which Edwin AI is applied to the various tasks facing hybrid operations and to add the capability to identify and trigger automated corrective responses. LogicMonitor is not rushing to push these out, however. Rather, the company is taking a measured approach to ensure that such capabilities are effective, accurate, and capable of adapting to the unique characteristics of each customer environment.

Conclusion

Hybrid observability is fast becoming the minimum table stakes for operations teams working to support modern applications. AI can be a huge help in keeping up with data glut and complexity, but only if the AI engine can be fed a full, rich, and current set of data upon which to operate. LogicMonitor's LM Envision platform with Edwin AI may very well be the first solution that has figured this out.

About Enterprise Strategy Group
TechTarget's Enterprise Strategy Group provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com