bmc helix

E-book

Managing Business Service Performance Across Cloud, Containers, and DevOps with AIOps

New approaches to greater observability into complex, and fast changing enterprise environments

Introduction

IT operations teams face a critical challenge. As the pace of digital business accelerates and enterprise IT grows more complex, it becomes both more important and more difficult to aggressively manage business service performance, improve stability, and predict and prevent performance degradation and outages.

To remain competitive, companies need to deliver high-quality service with reliable performance and availability for business users and customers. This depends on the ability to detect anomalies, find root causes, identify trends, and predict problems across the enterprise environment. However, that environment now encompasses more places than ever, from the data center to the cloud, and changes constantly due to the rise of DevOps, CI/CD, and containerization. The volume of log data is exploding as well, drowning ITOps in raw data and signal noise. As traditional approaches to monitoring fail to keep pace with the diversity of systems and environments in place, organizations struggle to capture both known unknowns and unknown unknowns.

This scenario has driven the rapid adoption of artificial intelligence and machine learning (AI/ML)-powered IT operations, or AIOps. A new generation of tools and methods now let ITOps teams leverage AI/ML and predictive analytics for advanced anomaly detection, event noise reduction, topology-aware probable cause analysis, and root cause detection. As the environment becomes more distributed, including public cloud environments, AIOps aids observability by making it possible to infer the state of systems that can't be monitored directly. By sorting through and analyzing data faster than human operators can process, AIOps provides clear, actionable insights for a faster and more effective response to problems.

In this e-book, we'll explore the essentials of AIOps for business service performance management, including:

- How monitoring and observability work hand in hand
- Leveraging AIOps to connect observability silos for faster probable root cause isolation
- Gaining AI/ML-powered service insights
- The next generation of AIOps visualization
- Driving observability in transforming cloud and container infrastructure
- Multidimensional visualization for today's layered topology
- How to put these practices to work using BMC Helix tools
- BMC's position as an AIOps leader

How monitoring and observability work hand in hand for performance management

To optimize system availability, ITOps teams need to closely observe and monitor metrics and datasets related to service performance—especially during upgrades and code launches. However, as companies adopt microservices and containers to meet the demand for more features, faster, business services are increasingly consumed as distributed functions across multiple layers of infrastructure and platform services. As a result, interdependencies have become more and more fragmented, making it hard to visualize the full IT stack. Assets in public cloud environments pose a particular challenge.

These trends are driving an urgent need to enhance monitoring and observability. While both provide visibility by bringing system data into dashboards, these are distinct and separate capabilities—and both have a critical role to play for enhanced end-to-end visibility.

Monitoring: Collecting raw, component-level data—when possible

A foundational element of traditional IT operations, monitoring consists of instrumenting specific components of infrastructure and applications to collect data on external behavior. By interpreting metrics, events, logs, and traces against thresholds, known patterns, and error conditions, IT teams can gain meaningful and actionable insights. Relying on this comparison of norms with anomalies, monitoring is most effective in relatively stable environments such as an organization's own data center, where key performance data and normal behavior are known factors.

Now cloud, DevOps, and exploding data volumes, including mobile devices and IoT, have created a situation where monitoring is no longer an effective approach. For example, consider a large, complex business service made up of multiple applications spanning public and private clouds, a diversity of distributed infrastructure, shadow IT assets, and perhaps a mainframe on the back end. It's just not practical to try to understand the health of such a service by selecting specific components to instrument for telemetry, and then keeping an eye out for threshold breaches and events. Lacking complete knowledge of key performance data and error conditions across so many systems, some not directly accessible, ITOps will end up with a mountain of uncontextualized data, unnecessary alerts, and false flags that slow issue resolution.

Observability: Inferring the unknowns from the knowns

Observability goes beyond collecting raw data to infer the internal state of a system from knowledge of its external outputs. Simply put, it's a method to learn about what you don't know from what you do know. This makes it well suited for modern enterprise IT—but it can still be highly challenging due to the volume, variety, and velocity of external data to be gathered, as well as the computational power and domain knowledge needed to make sense of it in real time.

For this reason, observability and AIOps go hand in hand. With the help of AI/ML, an AIOps tool can consume all available data from all system, aggregate it into a high-performance data store, and combine it with a complete topology of assets, systems, and applications to build a comprehensive model of relationships and dependencies.

Build on an open and integrated data lake ingested from every type of source—on-premises, private cloud, and public cloud—this modern observability platform approach makes it possible to produce actionable insights across the hybrid environment. On this foundational layer, domain-informed AI/ML algorithms can be applied to determine which externally observable data are correlated with which services, and infer the health of those services from their behavior.

To learn more, including the key role of observability in DevOps:

+ Read the blog on Observability vs Monitoring: What's The Difference?

AIOps connects observability silos for faster probable root cause isolation

Modern distributed systems and sophisticated applications keep people connected better than ever before, but our dependency on the digital platforms that make it possible puts immense pressure on IT organizations to avoid downtime and prevent outages across incredibly complex environments.

Impact analysis for services spanning application to network to cloud to mainframe has known gaps that can significantly affect service availability if they're not addressed. Until now, these gaps would require human intervention and never-ending bridge calls, with siloed teams racing to improve the mean time to repair (MTTR) score, an inefficient process that diminishes the customer experience and tarnishes the brand.

Part of the challenge is the wide spectrum of monitoring tools across the many layers that make up a digital service because each generates its own logs or metrics to understand performance. This wide spectrum of monitoring techniques eventually generates key artifacts such as metrics, alerts, events, logs, and topology that are unique and useful in the given solution but operate in silos and do not provide end-to-end service impact analysis.

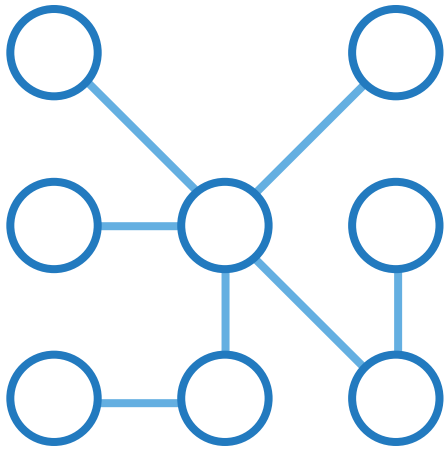

Finding a needle in a haystack

With all these tools, IT managers across teams experience a lot of "noise" from all the data and events generated, which makes distinguishing meaningful alerts from that noise extremely difficult. Each layer that contributes to a business service must be monitored and understood to deliver exceptional customer experiences.

Now, with innovations around ChatGPT-style algorithms, IT organizations can significantly improve the triage process and user experience. However, for those algorithms to deliver accurate, positive results with a high level of confidence, good data ingestion is required. Organizations rely on monitoring tools to gather data from logs, topology, metrics, alerts, and events. With metrics and logs, it is possible to create meaningful events based on anomaly detection and advanced log processing.

AIOps solutions can help IT managers avoid downtime with predictive, proactive insights across the entire application structure, from end user to cloud to data center to mainframe—and more quickly recover with accurate root cause identification when an outage does occur. AIOps removes the need for time-intensive investigation and guesswork so your team can see and respond to issues before they affect the business and shift their focus to higher-value projects.

To learn more about how AIOps connects observability silos for faster root cause isolation:

+ Read the blog on How BMC HelixGPT-Powered AIOps Connects Observability Silos for Faster Probable Root Cause Isolation

Gaining AI/ML-powered service insights

The more systems you have, the harder it becomes to keep watch over them all. To meet business and customer expectations, you need to be able to anticipate, avert, and troubleshoot incidents as fast as possible—but root causes can come from unlikely sources. Given a dynamic infrastructure with tens of thousands of hosts, containers, and services, you can't always anticipate where an issue would originate or what impact it would have on your organization. Anomaly detection, outlier detection, and composite alerting can reliably alert on the issues, but it's often hard to identify specific patterns to account for an incident. Meanwhile, incidents in areas where you haven't set alerts, such as an increase in latency or a spike in error rates, can result in significant service unavailability.

Continuous application monitoring and data analysis

The BMC Helix Operations Management with AIOps Service Insights feature helps ITOps make sense of overwhelming data and precisely identify hard-to-pinpoint trends. Build on established BMC ML features such as anomaly detection, noise reduction, and predictions, Service Insights uses an AI/ML auto-detection engine to provide a health score for each business service, uncover its associated events and anomalies, and provide root cause isolation and analysis.

When Service Insights detects a pattern or trend, the solution provides a plain-language summary of what happened, including the state or severity of a service, how long it has been that way, whether its health has improved or degraded over a given period, and if it needs immediate attention. ML pattern recognition shows ServiceOps engineers the precise moment when service performance degraded so they can take fast action.

Service blueprints

Showing a complete, end-to-end picture of how a business service is delivered, a service blueprint can be a powerful tool to understand both its user experience and its underlying components. BMC Helix Operations Management with AIOps provides a set of default Service Blueprints as well as simple ways to define your own. Once built, the solution can use traversal rules to find matching related nodes and include them in the blueprint as well. As the data changes, the system automatically re-applies these rules so that the service model stays up to date.

Natural language situation summaries

The BMC Helix Operations Management with AIOps Situations feature offers insight into events associated with a service based on based on factors such as occurrence, message, topology, temporal relationship, or a combination of these from across infrastructure, application, and network. An AI/ML-based event processing technique makes it possible to identify event patterns from hundreds of raw events, filter out noisy events, and automatically groups similar events together. A "Situation Summary" gives a human-readable insight based on natural language processing to describe the problem, why it occurred, and if it needs immediate action based on the underlying cause and severity of the problem.

To learn more about AI/ML-powered service insights:

+ Read the blog on Service Insights Powered by AI/ML

The next generation of AIOps visualization

Much of the attention around AIOps focuses on using AI/ML to analyze data, identify issues, and act automatically to remediate them in real time. But what about situations where ITOps or site reliability engineers (SREs) need to intervene?

When humans need to understand information quickly and clearly, the way it's presented is critical. Many aspects of information presentation have recently been undergoing an evolution, such as single-metric views giving way to consolidated dashboards, the overlay of events on metrics, service health scoring, and timeline views that correlate metrics, events, changes, and incidents. Now it's time for topology visualization and service modeling to evolve as well.

Three-dimensional visualization

Going beyond shapes, colors, and images, ITOps teams can now access three-dimensional models to represent different layers of a technology stack such as containers, database, network, compute, and so on. A layered topology view separates logical groupings of technologies into layers that can be visualized in relationship to each other or abstracted from the full topology view.

This three-dimensional visualization is simple, intuitive, and easy to navigate. Able to see immediately both the location and depth of the issue, IT operators or SREs can bring in only the most relevant teams to work on resolution. Provided with important contextual information that is clearly presented, they can then investigate the issue quickly, easily, and accurately.

To learn more about AIOps visualization:

+ Read the blog on **Power of AIOps Visualization: Moving from One Dimension to a Layered Topology**
+ View the two-minute video demonstrating layered topology in our **BMC Helix Operations Management with AIOps solution.**

Gaining full AWS visibility while controlling costs

Insights from logging can provide a great deal of context around the behavior of your applications and services, and help you troubleshoot when you're dealing with an outage. However, the more applications and services you run in AWS, the more complex your logging needs become—including end-to-end visibility into applications running on AWS services as well as into the AWS services themselves. AWS provides CloudWatch to centralize and aggregate all your log data, combined with data from other sources across your hybrid and multi-cloud environment. BMC Helix Log Analytics helps you store, monitor, and analyze this data.

BMC Helix Log Analytics makes it easier for ITOps to derive meaningful insights from logs. Used with monitoring metrics data from AWS and BMC Helix Operations Management with AIOps, Log Analytics provides fully contextualized data about the state of your AWS services and the applications running in it. Filtering by log metadata lets you zero in on specific components. Logs enriched with relevant data deepen understanding. Customizing alert messaging triggered by specific conditions occurring in the logs, combined with automated alerting via ML-assisted anomaly detection, help teams detect and respond quickly to developing issues.

As logs are collected and stored, ITOps teams can search, discover, or query any log record for further analysis. Data can be visualized in the form of a line chart, stacked chart, or numerical metric. Dashboards can be used to track the most important metrics, with alarms added to widgets for quick and simple monitoring.

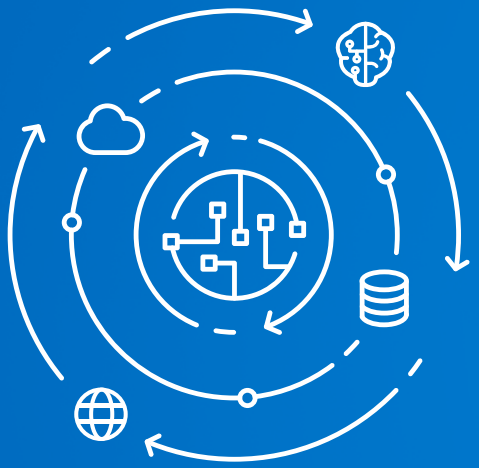

Archiving for compliance, analytics, and cost efficiency

Archiving is an essential part of visibility. Regulatory data retention policies can vary from several months to several years; for example, section 802 of the Sarbanes-Oxley Act (SOX) requires organizations to archive their data for at least seven years. Archived logs also make it easier to identify patterns and trends over a longer period than rolling log files, especially as older log entries get overwritten, deleted, or lost. Historical analysis helps teams run security investigations across large environments, conduct compliance audits, and perform long-term analytics on high cardinality datasets such as users, IP addresses, device IDs, and customer purchases.

As organizations seek opportunities to trim expenses, Log Analytics lets you archive and retain logs for a longer period of time at a cheaper cost. With the value of log data transitioning from high to historical in a matter of weeks or days, storage costs can quickly outweigh its potential value as a source of business insights. Compressing and storing archived logs in a location that does not need to be optimized for quick access can greatly reduce storage space and expense. In the event archived log data is needed for on-demand troubleshooting or any other operation, it can be decompressed and loaded into active databases at any time without any data loss.

To learn more about Log Analytics:

+ Read the blog on AWS Cloud Observability with Log Analytics
+ Access BMC Helix Log Analytics product documentation and video

To learn more about log archival and restore:

+ Read the blog on Archive logs to optimize storage & gain full visibility

Monitoring dynamic infrastructure

Many modern enterprises are now adopting a containerized, microservices-based architecture, but most monitoring solutions and data collectors still rely on a server-based monolithic architecture unsuitable for these dynamic environments.

The advantages of containerization are clear. Packaged with all of its dependencies into a portable container, software code can run uniformly and consistently on any infrastructure. The ability to start, create, replicate, or destroy containers in seconds allows tremendous agility. Container isolation and resource sharing enhance resource efficiency and density, helping optimize costs. High scalability ensures consistent availability and performance for fast-growing services. But to realize these benefits, ITOps has to be able to maintain clear visibility into their constantly-growing container environment.

BMC Helix Monitor Agent–Containerized meets this need by helping companies keep pace with every change in their dynamic infrastructure. Any change in your environment triggers an immediate update to your service model, helping you stay current with the latest data from your tools with no added effort.

Delivered in a lightweight, containerized collector, BMC Helix Monitor Agent–Containerized is pre-configured and ready to be deployed as a unified agent for vendor-agnostic public cloud, infrastructure, and database monitoring. Knowledge modules installed in the container support a broad range of monitoring types. Alarms can be set based on out-of-the-box BMC Helix Operations Management policies. Built for high scalability, reliability, and enterprise-grade performance, BMC Helix Monitor Agent–Containerized updates monitoring for modern dynamic architecture.

To learn more about BMC Helix Monitor Agent–Containerized read the blog:

+ BMC Helix Containerized Monitor Agent: Next Wave of Monitoring Evolution

To learn more about containers, read the blogs:

+ Containers Fit in DevOps Delivery Pipelines
+ How Containers & Kubernetes Work Together
+ Containers as a Service (CaaS) Explained

Extending monitoring and observability to Kubernetes

The rapid adoption of microservices, distributed applications, and containers has vastly expanded the challenge of monitoring and observability. While Kubernetes can simplify the management of containerized applications and services across different cloud services, it's also highly distributed, dynamic, and complex, introducing new layers and abstractions that expand the components and services to be monitored. To see whether your system is operating as expected, and be alerted when it isn't, you need to be able to take a proactive approach to Kubernetes with an observability strategy to keep track of all the dynamic components. By monitoring and debugging at the container, node, and cluster level, you can then drill down as needed for troubleshooting and incident investigation, and to view trends over time.

BMC Helix solutions solve Kubernetes visibility problems. For monitoring, BMC Helix Operations Management provides the ability to collect metrics from Kubernetes. For observability, BMC Helix Log Analytics allows you to track errors, monitor the health of containers that host applications, and even fine-tune the performance of containers.

Kubernetes logging with BMC Helix Log Analytics

Rather than trying to collect every log from across your pods and clusters, the BMC Helix Log Analytics Kubernetes integration automatically collects Kubernetes logs for you, regardless of the format they are written in or where in your Kubernetes environment they're stored. By automating Kubernetes log collection and analysis, you can avoid being overwhelmed by the complexity of Kubernetes logs. Rather than struggling to figure out where each log is stored and how to collect it before it disappears, you can focus immediately on actionable visibility.

BMC Helix Log Analytics integrates with BMC Helix Operations Management so you can analyze Kubernetes log data alongside metrics and other crucial sources of Kubernetes visibility. This provides full observability and contextualized data about the state of your cluster and the applications running in it. With a seamless and streamlined workflow for IT monitoring, troubleshooting, and investigation, SREs, DevOps engineers, and developers can easily go from problem detection to resolution in minutes.

To learn more about BMC Helix Log Analytics capabilities:

+ Watch the BMC Helix Log Analytics **overview video**
+ Access BMC Helix Log Analytics **product documentation**

Discover the AIOps toolkit designed to supercharge your business outcomes

Implementing AIOps solutions into your business will provide you with actionable intelligence for optimizing performance, giving DevOps, SREs, and ITOps increased agility by staying ahead of any potential service degradation or outages.

Leveraging a platform built with AIOps at its core will help increase efficiency, productivity, and innovation with your business.

To see AIOps in action check out our BMC Helix Operations Management with AIOps solution; sign up for a free Guided Demo **here**.

Forrester names BMC a leader in AIOps

Forrester recognized BMC in its *The Forrester Wave™: Process-Centric AI for IT Operations (AIOps), Q2 2023* report, an achievement that, in our opinion, validates BMC's commitment to innovation, customer satisfaction, and our continuous pursuit of excellence. Our cutting-edge, process-centric AIOps solutions represent the future of IT operations management, in which traditional approaches are enhanced and augmented with AI-driven intelligence.

To learn more about why BMC is a recognized leader in process-centric AIOps:

+ Read the blog on BMC Named a Leader in Forrester Process-Centric AI for IT Operations (AIOps) Wave

About BMC

BMC works with 86% of the Forbes Global 50 and customers and partners around the world to create their future. With our history of innovation, industry-leading automation, operations, and service management solutions, combined with unmatched flexibility, we help organizations free up time and space to become an Autonomous Digital Enterprise that conquers the opportunities ahead.

BMC—Run and Reinvent

www.bmc.com