

Ready for Business: A Modern Mainframe Requires Intelligent Management

Today's mainframe must be always on and always available. This requires a simplified, intelligent, and fast approach to finding and fixing potential issues.





Cloud computing, mobile technology, and the Internet have been adopted to shape responses to meeting customer needs.

The mainframe has been an IT computing staple

for decades. These workhorses do the heavy-lift processing for mission-critical applications across industries — and will continue to do so, according to mainframe market-growth forecasts from **Allied Market Research**.

That said, organizations are starting to realize they must both modernize mainframe technology and simplify its management. Thanks to the Internet, business always is open. This means the mainframe always must be available for data processing. Downtime is not acceptable.

Unfortunately, when one component of the mainframe goes offline, it's a big problem. It causes a tailspin among IT staff to find and resolve the issue. Simultaneously, customers get frustrated, and the longer they wait, the greater the likelihood they'll take their business elsewhere.

This Tech Dossier examines how artificial intelligence (AI), machine learning (ML), and automation are coming together to predictively detect potential mainframe issues before they become problems. This AI-driven operations (AIOps) approach delivers reliable, actionable insights to IT mainframe teams for quick resolution — ensuring customers and the business are happy, while IT staff can focus on strategic initiatives.

The Challenges of Supporting Today's Mainframe

Customers expect an always-on experience. This demand for always-available data and applications has caused enterprises to evolve their IT operating environments. Cloud computing, mobile technology, and the Internet have been adopted to shape responses to meeting customer needs.

At the same time, these trends have developed into a perfect storm for IT organizations. Staff now are trying to support and maintain a complex IT environment that operates across on-premises and cloud environments, as well as distributed geographies. This IT environment will continue to evolve as organizations create, ingest, and process ever-increasing volumes of data — much of which will be fed into the mainframe for processing.

The mainframe already is integrating with multiple data sources and applications, making it difficult to identify and source problems. Today's IT staff typically are monitoring by alarm and exception, whereby the problem isn't identified until an alarm is raised. For example, a cloud application can cause a mainframe outage, yet it may go undetected until frustrated customers report issues. There never is a good time for these incidents to occur. If they happen on the weekend or overnight, they may not be noticed until business hours resume; if the issue occurs during peak transaction processing time, numerous customers may be affected, further compounding the situation.

Many systems are configured to recognize "known" issues, missing exceptions to the "rules." Until a problem is identified for the first time, it may go undetected until the alarm is raised.

Furthermore, reactive response time to alerts is slowed by a dearth of mainframe skillsets, which rarely are taught in today's university curriculum.

Network World has been writing about the problem for several years now, **explaining**: "People who started their IT careers in the 1970s and 1980s — when the mainframe was king — are now Baby Boomers at the end of their careers. The generations behind them took up different computing platforms, meaning there are few people to pass the mainframe torch to."

Because of this gap, few enterprises have the ability to effectively and efficiently marry Big Iron skills with modern systems.

"For example, there is a lack of IT talent who can both support and manage the mainframe while integrating it with cloud applications," said Alan Warhurst, Senior Manager of Product Management at BMC Software, who also works with the BMC Mainframe Executive Council.

Warhurst added, "It's becoming increasingly critical to have capabilities that align with the needs of a modern mainframe."

Enterprises must keep pace with the evolving IT landscape by ensuring the mainframe can adapt to whatever comes next. This means moving beyond just monitoring for problems, and being able to predict and resolve issues at speed. To do so, companies need machine power and intelligent systems to aid existing IT teams.



DEFINITION

AIOps



Multi-layered technology platforms that automate and enhance IT operations through analytics and machine learning. AIOps platforms collect a variety of data from various IT operations tools and devices to automatically spot and react to issues in real time while providing traditional historical analytics.

The Opportunities in Adopting an AIOps Approach

AIOps can help address those operational and skillset challenges associated with the mainframe.

AIOps presents opportunities to better manage and simplify the mainframe environment, delivering benefits to both IT and the business. For example, intelligent systems reduce downtime. AIOps uses machine learning and automation to learn what "normal" is, then looks across the mainframe to detect issues, rather than waiting for an incident or problem to occur. An AIOps approach also provides predictive, actionable insights — contextual information that mainframe teams can use to speed resolution.

Ultimately, this intelligence accelerates business value. Instead of reactively fighting fires, IT staff are freed to focus on strategic efforts. Applications are always available, ensuring excellent customer experiences, and the right AIOps solution reduces the total cost of ownership (TCO) of mainframe management by overcoming the skills gap, and simplifying issue detection and response.



AIOps offers businesses the ability to understand and address mainframe complexity amid an evolving IT operations environment.

As companies explore the benefits of AIOps, they're increasingly showing interest in this approach. In a recent **webinar poll**, BMC Software found that 70% of companies are looking at potential-use cases, and 21% are formulating deployment plans. These results highlight that organizations require a realistic approach to AIOps adoption, with practical considerations.

Best Practices When Considering AIOps Adoption

AIOps offers businesses the ability to understand and address mainframe complexity amid an evolving IT operations environment. However, it should be more than just a tool for the toolbox. It also must deliver business value. Below are three considerations with associated best practices as you explore AIOps solutions for the mainframe.

1. Cost of ownership

Ultimately, AIOps should reduce TCO. That means taking a hard look at existing in-house skillsets. Consider how soon you'll need to refresh mainframe skills, whether those skillsets will be available, and at what cost. This includes support and maintenance, as well as the data scientists and domain experts who rely on the mainframe's data-processing efforts. These talents come together to provide context and get the most from the mainframe.

There are additional (and sometimes more significant) cost considerations pertaining to hardware and software. Understand the costs for the mainframe to process data from different applications, and the speed at which it's capable. This ties into the prerequisite hardware and software resources for data processing. Do you need to build a data center just to keep up with the sheer volume of data hitting the mainframe?

Best practice: As you research AIOps solutions, take cost factors into account and ask: Does the solution help solve the need for skills and expertise? Does it address today's mainframe processing capabilities and the expectations for future needs?

Detect Mainframe Issues Before They Become a Problem

BMC Automated Mainframe Intelligence (AMI) Operational Insight reduces downtime and simplifies mainframe management by leveraging an AIOps approach. The solution is focused on mainframe availability and fast detection of problems. It provides actionable insights so teams can take action before systems are affected.

BMC AMI Ops Insight delivers critical benefits, including:

Immediate value: BMC AMI Ops Insight is the application of expertise, not just tools. It offers built-in data science and domain expertise to alleviate heavy lifting among mainframe teams. The solution optimizes the data fed into the mainframe, and how it is interpreted, thus reducing the need for in-house staff to perform analysis and determine appropriate metrics. It also reduces the overhead needed to manage the mainframe; with improved detection analytics and faster issue resolution, the mainframe runs optimally.

Speed and accuracy: BMC AMI Ops Insight predicts potential problems to reduce and prevent performance issues. Using multivariate analysis and machine learning, it detects anomalies before they impact the mainframe, while reducing false positives. It also decreases MTTR by delivering predictive, actionable insights about potential issues.



A modern mainframe: By simplifying management and support, BMC AMI Ops Insight ensures the mainframe is ready to handle increasing data volumes, while meeting the needs of the evolving IT environment. The solution is always learning and getting smarter to best support the mainframe.

“The mainframe shouldn’t be treated as a silo,” said Warhurst. “Our AIOps approach, combined with our mainframe expertise, simplifies mainframe management to empower IT teams while ensuring the business is always on and always available.”

For more information about BMC AMI Ops Insight, visit: <https://www.bmc.com/it-solutions/bmc-ami-operational-insight.html>

2. Know what you’re solving

Behind the exploration of potential-use cases for AIOps is the need to identify the problem and urgency for resolution. For example, if speed of detection and resolution might be the issues, get some facts. Answer the following questions:

- **If your mainframe goes down,** how fast can the problem be detected and resolved?
- **What are the ramifications** of being offline?
- **What exactly is slowing you down?** Is it tool capabilities, a lack of staff, or insufficient mainframe integration with applications?
- **What does the business expect** in terms of mean-time-to-resolution (MTTR), and what are you achieving now?

In addition, determine if problem detection and resolution are accurate. For example, find out how much time mainframe teams spend chasing false positives, then consider the time it will take for existing staff to train ML models to look across all KPIs to ensure nothing is missed.

Best practice: Ultimately, detection and resolution revolve around speed and accuracy. No amount of mainframe downtime is acceptable. Ask vendors how their AIOps solutions use automation and ML for faster identification of potential problems. Also, find out how much training your existing mainframe team will require to get up to speed for the solution.



An AIOps approach takes into account people, processes, and technology.

3. Prepare your mainframe for the future

Just like the IT environment itself, the mainframe must evolve. It must be ready for new devices and applications, continual data generation and data processing, as well as the ongoing mainframe skills gap.

An AIOps approach takes into account people, processes, and technology. It helps IT stay ahead of potential disruptions, reducing the burden on stretched-thin teams, while speeding detection and response processes. In addition, the combination of technologies such as ML, data science, and domain expertise in one solution simplifies management and support. This ensures the mainframe is ready to handle increasing data volumes.

Best practice: The modern mainframe supports the ever-evolving IT environment, using an intelligent management approach. Yet to future-proof the mainframe, it's critical to go beyond off-the-shelf AI and ML tools. Seek a partner with the domain expertise to optimize the mainframe and make it work for you.

The Bottom Line

The mainframe is and will continue to be a critical component of IT services. Both internal and external customers rely on its data-processing efforts. Simply monitoring for problems isn't enough; companies must quickly detect issues before the alarm is raised and the mainframe suffers downtime.

An AIOps-driven approach to predictive detection ensures there is little to no business disruption, while simplifying mainframe management.

"AIOps finds and fixes problems quickly," said Warhurst. "This not only simplifies management for mainframe teams, it also helps them deliver strategic value to the business."

Ensure your mainframe stays ahead of business disruptions. Get more information at:

<https://www.bmc.com/it-solutions/bmc-ami-operational-insight.html>



Can you put your trust in Alops?

As ops tools morph into Alops, they bring greater value to more complex deployments—if users will let them.

By David Linthicum

Alops (artificial intelligence for IT operations) is one of those cool buzzwords that is actually part of another buzzword: cloudops (cloud operations), which is a part of the mother of all buzzwords: cloud computing.

The concept of Alops and the tool category of Alops are really the maturation of operational tools in general. Most of those in the traditional ops tools space, at least in the past few years, bolted an AI engine onto a tool and called it Alops.

Some purpose-built Alops tool startups out there are leveraging AI from the jump. All are worth a look as you select Alops tools; however, there are no mainstream brands.

The objective was and is obvious. Since most of these tools have been data gathering tools and analytics tools from the beginning, adding AI allows them to learn from that data rather than just externalize issues with the services under management. In some cases, they can correct issues using preprogrammed routines, such as restarting a server or blocking an IP address that seems to be attacking one of your servers.

Now that we're a few years into this paradigm and its technology offerings, we're starting to note some patterns—some good, and some not so good. Let's explore both.

As far as what's working, Alops tools in many instances are ops tools in their fourth, fifth, or sixth generations. Moreover, most of them have had public cloud management in mind for a while and are able to bridge the gap between on-premises legacy system management and managing applications and services in the public clouds.

They are capable tools for managing and monitoring cloud, multicloud, legacy, and even IoT and edge-based systems. This ability to support complex system heterogeneity is really the true value of the ops tools, and why they are important to those implementing cloud or noncloud systems.

The downside seems to be that most users are not taking advantage of the AI subsystems in the tools, so you may be paying for a feature you're not using. I don't view this as the fault of the tool provider; for the most part, this relates to how the tools are installed, set up, and used by the traditional and cloudops teams. This is due to a lack of training in some cases, or a lack of valid use cases in the current set of systems—cloud and not cloud—under management.

To read the rest, [click here](#)



How AIOps improves application monitoring

Devops and site reliability engineers are vital to keep applications functioning. AIOps boosts effectiveness another notch.

By Isaac Sacolick

IT operation teams use many tools to monitor, diagnose, and resolve system and application performance issues. In a recent survey of 1,300 IT professionals on the future of monitoring and AIOps, 42 percent report using more than 10 monitoring tools; 19 percent use more than 25 tools.

That's a lot of technology just to keep the lights on and provide the data required to monitor, alert, research, and resolve application incidents.

Monitoring tools are not one size fits all, especially for organizations running mission-critical applications in multi-cloud environments. As organizations invest in mobile apps, microservices, dataops, and data science programs, new monitoring tools are being added to provide domain-specific monitoring capabilities.

AIOps platforms aim to simplify this landscape of monitoring tools. AIOps helps organizations that require high application service levels better manage the complexity of their monitoring tools and IT operational workflows. As the name suggests, AIOps brings machine learning and automation capabilities to the IT operations domain. These technologies aim to resolve incidents faster, identify operational trends that impact performance, and simplify the procedures required to resolve issues.

AIOps is an emerging platform. In the survey, 42 percent of respondents either had never heard of AIOps or had thought that applying machine learning to operations was “not a thing.” Only 4 percent are using an AIOps tool in production today. Although AIOps is an emerging platform, there's a solid business case for many organizations to consider it.

AIOps is driven by business need and operational complexity

More businesses today rely on applications to serve customers and run operations. That drives higher requirements and expectations on the reliability, performance, and security of the applications.

It also fuels demand for application development teams to build new applications and enhance them more frequently. The job responsibility of maintaining application service levels has also broadened during the past decade.

Once upon a time, organizations staffed the NOC (network operations center) as the front line of defense. If you ever walked into a NOC, you would likely see dozens of computer monitors with warning lights and trend visuals to help the staff pinpoint issues—ideally before an end-user experienced one and opened tickets.

Business and IT leaders began changing this model by introducing devops practices and site reliability engineers. Devops changes the IT department's culture by establishing a collective responsibility to enable frequent deployments and better support customer and employee needs. Tools and practices such as CI/CD (continuous integration and continuous delivery) and IaC (infrastructure as code) are part of what enables more frequent deployments.

But devops practices also require a shared operational responsibility ensuring that applications are reliable, perform well, and are secure. That means more people in the IT organization need access to all the different monitoring tools.

Many IT organizations also hire SREs (site reliability engineers) to connect development and operations. SREs take a software engineering approach to system administration topics. In another survey that targeted SREs, they indicate that incident response is a massive part of their job: 49 percent claim to respond to at least one incident every week.

Maturing devops practices and hiring site reliability engineers is how a growing number of IT organizations are facing increasing operational challenges. But just expecting them to make sense of the dozens of monitoring tools being used is a recipe for poor performance.

AIOps platform capabilities and technical architecture

How can AIOps improve the status quo? AIOps platforms typically have the following architecture components and capabilities:

- **A central data platform** for aggregating raw logs and data from different monitoring tools.
- **Out-of-the-box integrations** with the most common log formats, monitoring tools, IT service management tools, agile development tools, and other collaboration platforms.
- **Machine learning capabilities** to help identify patterns in the aggregated data.
- **Consoles, dashboards, and analytics** to help IT operations see and manage multiple systems from a central interface.
- **Automation capabilities** that enable IT to communicate status, route issues, and autorespond to common problems.



What differentiates AIOps from other IT operational platforms is the ability to aggregate data easily, leverage machine learning to find problems, and use automation as a tool to resolve them. AIOps doesn't replace the existing monitoring tools. It integrates with them so that more people in the IT department have improved visibility to problems without the complexity of learning and using multiple monitoring tools.

Similarly, AIOps platforms typically don't replace existing IT service management, workflow, agile, and other communication tools. Instead, they are a central platform to interface with them while alerting and resolving an incident.

Monitoring mission-critical applications without AIOps

Imagine your e-commerce application experiences slow performance when users try to complete a purchase. The first indicator that starts to send out alerts is the shopping cart abandonment rate.

The e-commerce leader quickly opens a ticket about the issue in Cherwell's mobile interface, but the IT team has already been alerted to the problem. As more users try to make purchases, the underlying Web servers hang and database connections stay open. Alerts from DataDog report these issues, and Splunk reports Java exceptions in the e-commerce application's log files.

To read the rest, [click here](#)



Not all Alops tools are created equal

These tools can solve multicloud complexity issues, but they use very different approaches.

By David Linthicum

Alops is becoming the new norm for operational tools supporting cloud operations. This technology can be applied to all types of operational tasks, providing intelligent automation that learns as it solves operational issues.

These tools must carry out preprogrammed self-corrective processes, and the Alops tools' ability to learn during those processes creates a huge advantage. For instance, understanding that performance issues could be saturation caused by cyberattacks should kick off security processes to mount a defense. Or moving out of a performance threshold should automatically launch more resources to bring performance back to an acceptable range.

The number of things you can do with these tools increases each day, and it's likely to be standard equipment for those of you deploying and operating multicloud.

Most important, Alops tools can deal with thousands of data points and make correlations that most humans would not make. Moreover, as they correlate these data points the tool itself is smarter—it knows what the information actually means and how to assist the cloudops team.

The trouble is that many products in this space are actually old technology made new. We've been using operational tools for years. Those tools were redone to support public clouds; now they have been rebranded as Alops tools with some built-in AI capabilities.

The trouble with this type of evolution is that it's happening so fast, the tools are naturally going to take different approaches. Some are very data driven, capable of analyzing historical data; other focus on real-time monitoring.

Data-oriented tools look for patterns in the data—typically assisted by an AI engine—in order to find cause and effect. They get to the root cause of an issue without the cloudops staff having to cull through gobs of data.

Also, AI is leveraged in different ways. Some have pretrained AI systems within the tools, which means the tools comes with a predetermined amount of knowledge. Others focus on training from scratch. Each approach has advantages.

To read the rest, [click here](#)