




Intellyx™



Accelerating AIOps on the Mainframe

Best practices for achieving operational resiliency through AIOps strategies including IBM Z

Jason English

Principal Analyst, Intellyx

April 2021



AIOps means much more than the combination of ‘Artificial Intelligence’ with IT Operations. It’s not just a specific tool you can buy, nor is it a rote technique that guarantees a certain result.

AIOps is a fundamental capability cultivated by organizations that brings together people, processes and technology, intelligently leveraging a deep awareness and understanding of all of their systems and data to improve operational agility -- and most importantly – *operational resiliency*.

Operational resiliency is an essential property that shields an enterprise’s digital lifeline from the most disruptive forces it must face, and this resilience under pressure prepares the company’s entire digital estate to constantly spring back stronger in the face of change.

A comprehensive AIOps strategy should take into account all efforts toward operational resiliency -- from ensuring a flawless experience for new customer-facing application functionality, to maintaining the security and availability of the enterprise’s beating heart -- the mainframe.

In our previous Intellyx whitepaper “[Why You Should Do AIOps with the Mainframe](#),” my colleague Charles Araujo clearly lays out the inseparability of the mainframe from the rest of a business’s digital function:



“The mainframe must be central to an AIOps adoption effort not because it is somehow more important than the rest of the stack, but simply because it is an essential element of most business-critical workflows.”

Fortunately, a lot of innovation has been focused on making mainframes flexible to stress and responsive to change, and many leading organizations have successfully navigated this change. This paper will discuss strategies for achieving best-in-class operational resiliency by making the modern mainframe a central part of the AIOps transformation.



The Journey to AIOps

Like a composite bow made of multiple layers of wood and polymers with different tolerances and tension ratings, we want the combined operational strength and flexibility of the enterprise's composite systems to be greater than the sum of its parts.

No single part of your application stack should be allowed to become hardened or brittle. Therefore, we want to augment our own human ability to perceive how our systems are behaving with automation and machine intelligence, so they are less likely to snap under stress.

The tension on our systems is drawing ever tighter, as companies are inevitably transforming toward a hybrid IT future. This means bringing forward the entire technology estate of applications, data, and core systems to interact with cloud infrastructure and cloud-based services – thereby setting the need for greater transparency and correlation of Ops data across that landscape to maintain operational resiliency at a business process and workflow level.

For example, picture yourself as the CIO of a regional bank, trying to come to grips with modernization and new application development to meet changing customer demands, while still maintaining consistency for critical ongoing business operations.

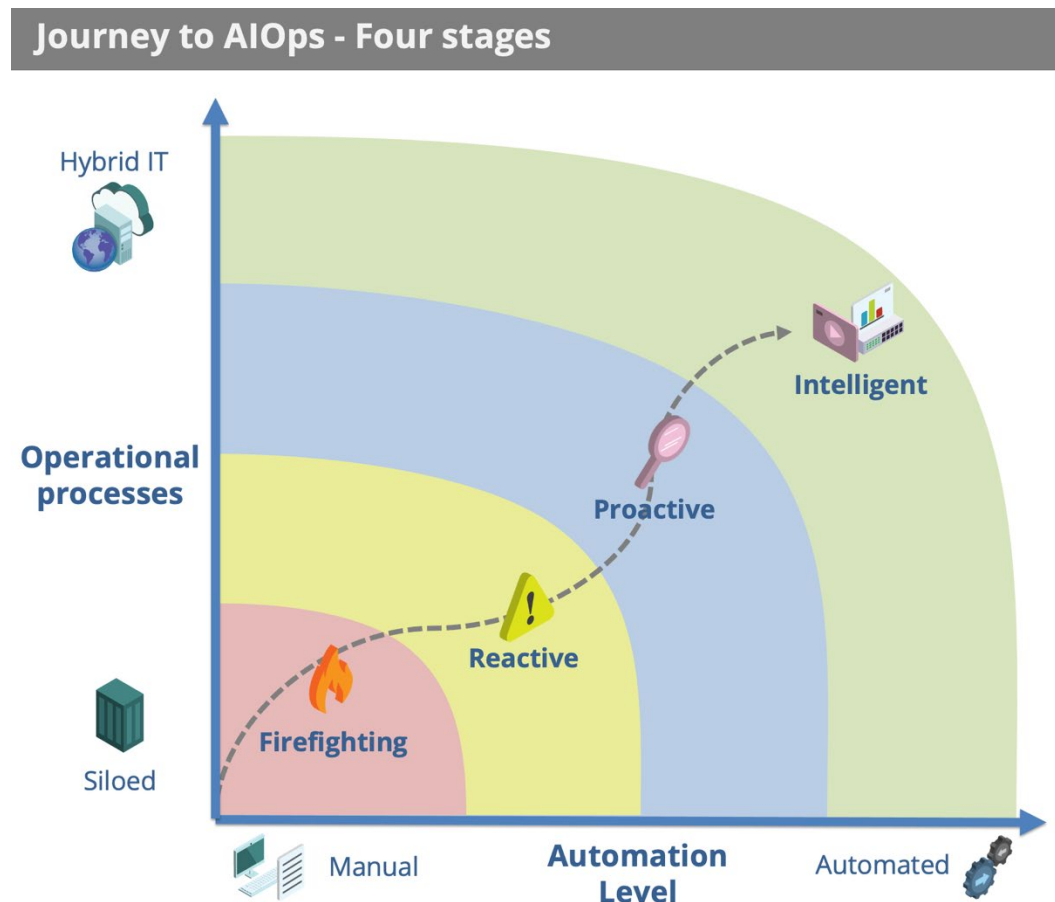
Everything is operating smoothly, availability is high, and teams respond to any issues that arise in minutes. But today's operational resiliency won't last forever. You are asked to deliver a new mobile deposit and payments app that integrates with your existing systems, without impacting your current customers.

With future business prospects depending on the new app, which itself must depend on a mainframe already soaking in current business requirements -- how would you plot your course toward AIOps resiliency, without breaking anything?



Stages of the Journey to AIOps

As part of a modern, high-scale, fast-changing business operation, you could plot your mainframe's course to AIOps resiliency along two dimensions.







[Figure 1. Journey to AIOps - Four stages of operation]

In the above chart, the 'x' axis goes from manual to fully automated IT operations processes, and the 'y' axis value ranges from siloed system-level processes to hybrid cloud processes.

Clearly, getting up and to the right is where we want to be on this chart. There would be some common experiences that an organization notices at each of these stages or modes of operation.



Stage	Experiences
Firefighting 	Bug hunting, constant paging, crisis management war rooms, break/fix sessions, late night/weekend work, nothing is repeatable
Reactive 	Monitoring for issues, trying for shorter defect identification and resolution times, sorting through logs, some root cause identification and documentation
Proactive 	Predicting failure scenarios, testing team and system response times, watching for change, performance and functional issues on mainframe and other services, some automated responses and playbooks
Intelligent 	Intelligent, policy-based operational management system, integrates distributed processes across enterprise, high automation, transparent observability, explainable AI, machine learning for filtered anomaly recognition

[Figure 2. Journey to AIOps – Experiences encountered at each stage]

Importantly, all the steps of the journey toward AIOps resiliency aren't exclusive to the mainframe platform itself. In fact, each improvement in maturity makes the mainframe a stronger, more transparent and more change-responsive member of the enterprise's entire application estate, from front-end apps and websites, through middle tier integrations and third-party service dependencies, all the way to processes and data residing on the mainframe.



Measuring improvements in practice

Time is money. Time lost due to system downtime, time spent recovering a lost order, time a customer might spend waiting on an unresponsive ordering app -- all of these come with a calculable burden of lost revenue and customer churn, along with the human toll of customer support hours and issue remediation labor costs.

Organizations generally target their AIOps goals and measure their performance by several 'mean time' metrics -- MTTD (mean time to detection) and MTTR (mean time to resolution) being the most common. In short, we want AIOps resiliency so the org can respond to change faster, and eventually automate away as many issues as possible.

Now that we've identified these modes of operations as milestones of our AIOps resiliency journey and learned how to recognize where we are on this course, let's focus on three primary practice areas of AIOps improvement involving the mainframe.

Practice 1: Monitoring and Observability

Goal: Use system-wide monitoring and observability to augment human and system visibility into the ongoing operational health of the entire application estate, and gain early awareness of emerging anomalies and resulting problems.

Issues: Gaps in coverage appear across interdependent systems, each with their own siloed monitoring tools. Operators spend time looking for problems rather than getting high-priority alerts or intuitive dashboards that make issues easy to sort out from the alert noise, resulting in missed incidents, and late reaction time.

Key Capabilities:

- Full-stack observability across hybrid IT applications and domain-specific mainframe workloads
- System-wide application performance monitoring (APM) for metrics, events and alerts, sortable at an application / domain / group level

AIOps KYA "Know Your Acronyms" Lingo

MTTD	Mean time to Detect
MTTI	Mean time to Identify
MTTK	Mean time to Know
MTTR	Mean time to Repair / Resolve



- AIOps identifies a baseline understanding of normal operations down to the mainframe level, for comparison and earliest possible detection of anomalies
- AIOps filters out alert noise and false positives, promoting significant issues and flagging them for human or automated intervention.

Value metrics: Faster MTTD (mean time to detect), fewer overall issues promoted as tickets, faster issue acceptance by teams, higher incident acceptance percentage rate. Some companies have reported an \$139K average labor cost savings per incident avoided.

Practice 2: Root Cause Analysis

Goal: Establish the root cause of incidents appearing underneath the UI or as alerts, isolating them within a highly complex system under current load conditions, and learn to predict similar issues in the future.

Issues: Diagnosing problems leads to dead ends and wasted time. Issues are much harder to diagnose and reproduce due to interdependencies and constant change in live systems. With a flood of data to sort through, it becomes hard to isolate and recognize the source of problems.

Key Capabilities:

- AIOps shines best at detecting and correlating code bottlenecks in metric data (logs, SMF data, etc.) in applications and at the z/OS subsystem level.
- Stream and correlate all log and SMF data to a single location for rapid analysis through historical and predictive lenses to spot emerging problems.
- Deliver results of investigation back to issue resolution platforms and/or incident resolution teams.
- Performance and availability analysis for capacity planning, modeling, and redundancy planning activities.

Value Metrics: Faster MTTI/MTTK starting points for issue resolution teams. Improved performance and reduced outages due to planning and preventative measures. Better SLO performance for customers, with reduced SLA failure financial risk. More accurate and right-sized capacity planning. Companies succeeding at this practice report up to 50 percent reduced labor cost of up-skilling IT operator team members with AIOps insights.



Practice 3: Intelligent Automation

Goal: Consistently resolve incidents with intelligent automation that reliably improves issue triage and resolution time, while augmenting the ability of IT operations teams to collaborate and act on high-risk and high-value remediation efforts that have the most impact on customer satisfaction and cost.

Issues: Constant firefighting, all-hands war rooms and long issue resolution times. Escalations to expert SREs and business domain SMEs are common as minor or preventable issues become mission-critical in production. Very little automation working in break/fix environments, poor team morale, high labor costs, and error-prone manual actions taken for each resolution cycle.





Key Capabilities:

- Reduced manual intervention, increasingly automated policies, documented response plays and automated runbooks for faster resolution times.
- Pinpointed code, component, workload and system-level context with insights to take action to reduce the blast radius of failures in multi-tiered environments.
- Improved collaboration on incidents with traceable event logs, reporting to issue resolution systems and ChatOps.
- Proactive resolution of incidents leveraging explainable AI to improve confidence in resolutions.

Value Metrics: Reduced MTTR for mainframe issues and system-wide problems. Reduced labor cost due to automated resolutions and remediation playbooks. Lower outage costs for SLA violations and lost revenue. Mid-to-large companies that rely on digital business have reported an average of \$420K saved for eliminating each hour of outage.



Clues for progressing on your journey

Stage / Activity	Monitoring & Observability (detection)	Root Cause Analysis (decision)	Intelligent Automation (action)
Firefighting 	Enable monitoring and automated alert reporting	Reduce trouble tickets and need for firefighting by pinpointing major issues among a storm of alerts	Eliminate costly issue resolution war room sessions and trial-and-error debugging
Reactive 	Use baselines and SLA boundaries with analytics to identify critical issues faster	Attain data, workload and application-level reporting, with machine learning to identify comparable issues, and better documentation	Filter and understand the SLO impact of events for triage and prioritized resolution
Proactive 	Aggregate all events in unified dashboards with intelligent reporting	Predictive identification and classification of emerging issues, automated cross-system reports and escalation workflows	Build cross-platform runbooks and response plays for faster resolution
Intelligent 	Explainable AI for transparently identifying and reporting issues that exceed SLO boundaries to the right team or service	Intelligent policy-driven management of known issue patterns and deep analysis of system-wide causes	Set policy-based automated resolutions using AI insights, and escalation of only mission-critical human remediation activities

[Figure 3. AIOps journey checklist for achieving operational resiliency by activity.]

Every company's journey to AIOps resiliency will be as unique as the needs of end customers and the configuration of the hybrid IT production environment your mainframe operates in.



As we follow our progress past each milestone of the journey, there are common areas for improvement that can be unlocked within each practice to continue moving forward to the next mode of AIOps automation and intelligence, with the right combination of research and enablement.

The Intellyx Take

Operational resiliency is a core success factor of digital transformation, and any organization with a history of meeting customer demand depends on the mainframe to support that change. Leave the mainframe out of the journey at your own peril.

Every significant business application we deliver must be architected to support a modern hybrid IT delivery environment, which by definition includes responsive user interfaces, microservices-based cloud resources, integration middleware, and the modern mainframe.

The status quo of telling IT Ops teams to chase down alerts once they make it into production won't cut it anymore. Companies must strive to become more responsive and intelligent -- **more resilient** -- to survive in an increasingly Hybrid IT landscape.

AIOps can help organizations get a clear view of what's going on, true, but there's also a lot of resilient capability for intelligent automation we can explore, once we lift the cover off this complex business engine.

Transparency and automation with intelligent insights can improve our strategy—for development and testing, provisioning, deployments, disaster recovery, cross-platform orchestration, workload scheduling—and ultimately, customer service objectives.



About the Author

Jason “JE” English ([@bluefug](#)) is Principal Analyst and CMO at [Intellyx](#), a boutique analyst firm covering digital transformation. His writing is focused on how agile collaboration between customers, partners and employees can accelerate innovation.



In addition to several leadership roles in supply chain, interactive and cloud computing companies, Jason led marketing efforts for the development, testing and virtualization software company ITKO, from its bootstrap startup days, through a successful acquisition by CA in 2011. JE co-authored the book [Service Virtualization: Reality is Overrated](#) to capture the then-novel practice of test environment simulation for Agile development, and more than 60 thousand copies are in circulation today.

Resources: Learn More

Visit the **AIOps on IBM Z** webpage for more information and assessment tools:
<https://www.ibm.com/it-infrastructure/z/capabilities/it-operations-management>

© 2021, [Intellyx](#), LLC. *Intellyx retains editorial control over the content of this whitepaper. At the time of publishing, IBM is an Intellyx customer. Image credits: Jason English, Intellyx (diagrams/cover art); Charles Araujo, courtesy.*