# Maturing in AIOps – creating value at each step of the journey

Enabling Government & Public Services Clients through the Implementation of AI-enabled Operations

# Contents



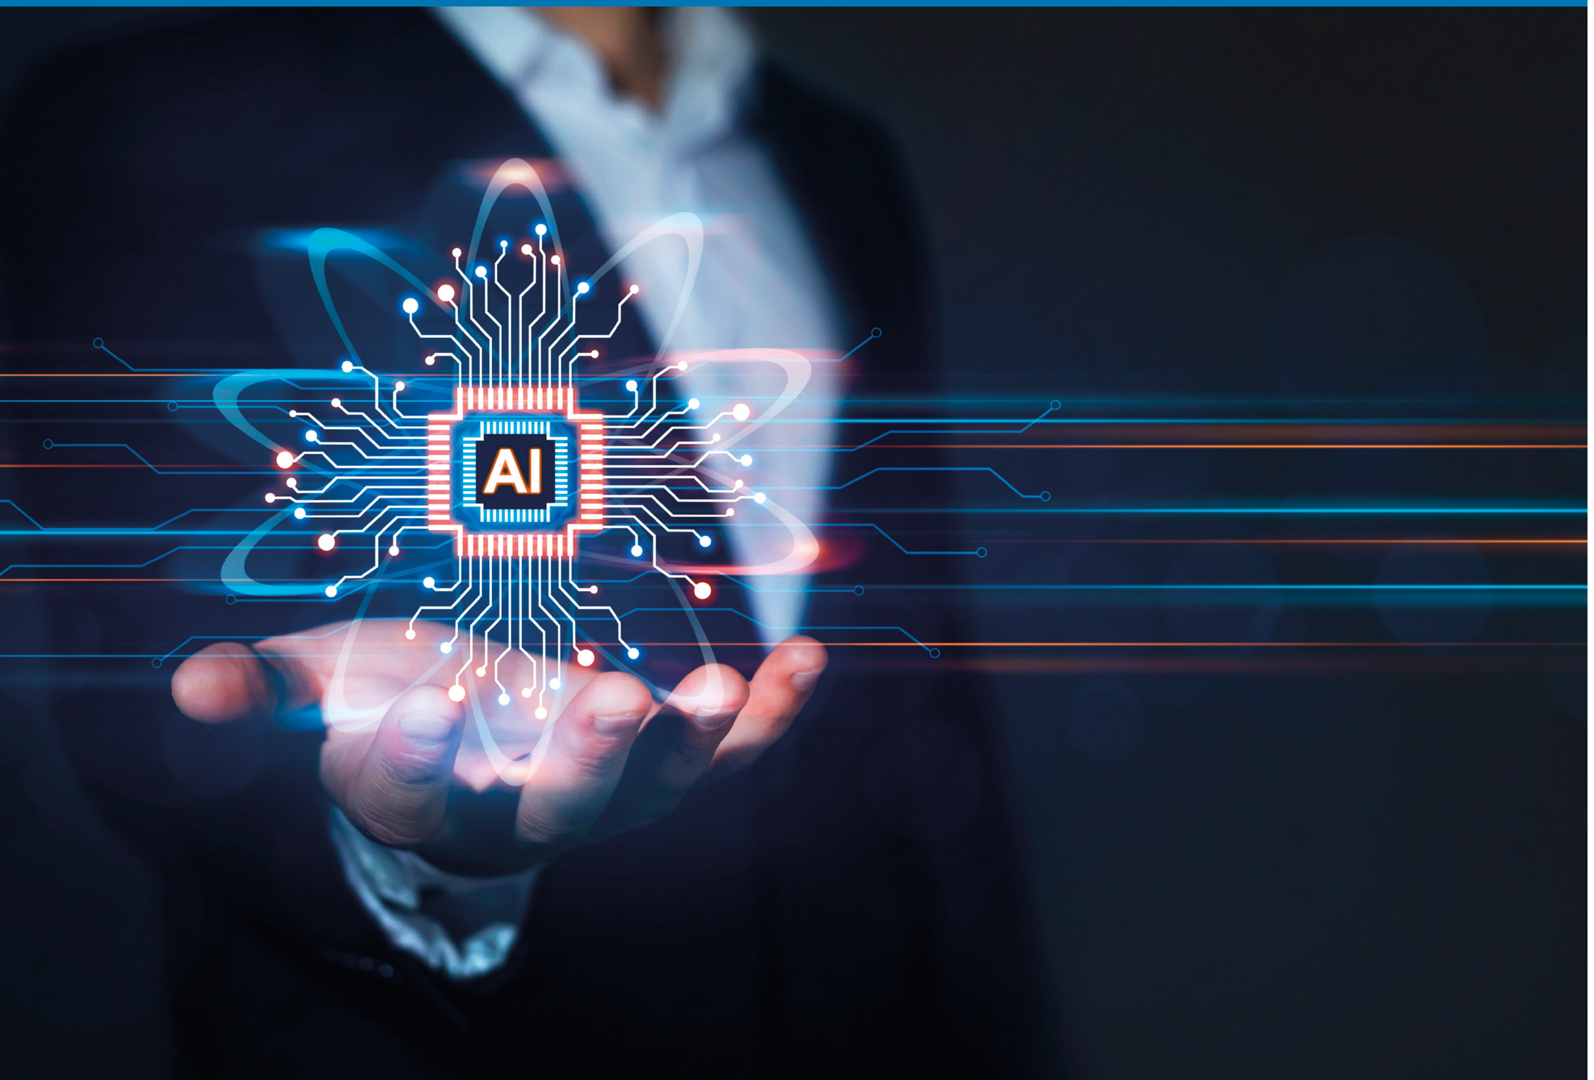


# Introduction

The world of science fiction has presented us with countless futuristic visions, from handheld devices that have the power to diagnose and heal, to computer systems that connect directly with our brains.

However, one concept that sticks out for many who lead IT operations is the idea that a starship's captain, can issue a simple verbal command for a technically complex action to execute instantly. Although exploring the various elements required to automate this process and technology is too arduous for a television series, examining the components of this "easy button" can help governments working to streamline their IT operations with the help of AI (Artificial Intelligence).

Artificial Intelligence for IT Operations commonly referred to as AIOps, has evolved from something of science fiction to an attainable reality that can improve operations and deliver results. Tool advances in observability, workflow automation, AI, and the connector libraries used to integrate them have created opportunities and new value for technical operations. Maturing the use of the tools individually can allow them to create efficiencies on their own; however, integrating them can scale those efficiencies and create new, stronger autonomous capabilities. If organizations seek the benefits of AIOps, IT leaders must mature the building blocks, demonstrate the ROI of maturing the components individually, and integrate where possible along the AIOps journey.

AIOps is the culmination of many efforts and initiatives intended to build workforce multipliers, reduce human error, and adopt an automate first disposition. IT leaders committed to the journey will likely benefit from maturing the component entities to improve quality, and manage ever increasingly complex IT environments.

This paper gives an overview of the journey and its potential benefits, introduces the components needed for AIOps and charts a course for getting started.

# Overview

## Fueling the journey by creating value in both the sum and the parts

The COVID-19 pandemic has been a catalyst for significant changes in the ways government agencies operate. One of the most notable changes has been the accelerated adoption of cloud technology. While this shift has brought many benefits, it has also introduced new challenges. To fully leverage the benefits of cloud technology, many agencies invested in DevSecOps initiatives that have helped integrate security into the development process from the outset, rather than adding it as an afterthought. However, this has also made IT more complex, as teams must now manage a range of different tools and processes to ensure that applications are secure and reliable. Another challenge that has arisen from cloud awdoption is the need to manage hybrid cloud environments. While many agencies have migrated to a mix of cloud providers to balance cost and performance and to avoid vendor lock-in, this can be difficult to sustain, as operations teams navigate a complex set of tools and processes.

With all this added complexity, CIOs and IT Leaders need solutions that better support their sustainment workforce in managing increased integrations and variables. Many of our clients turn to the modern delivery tools found in AIOps. The promise of self-healing, "touchless IT" where administrators login to one-to-many tools instead of servers, and full self-service provisioning scenarios is hard to ignore. But it is essential to recognize that achieving these scenarios requires a solid foundation built from the ground up.

To illustrate, let's consider a self-healing scenario. The first step is data collection from multiple observation points. This data is then analyzed by an observability platform, which creates a set of events. Machine learning algorithms curate the set to a single alert, which is then sent to a workflow tool that uses trained AI to determine the best course of action. Finally, an infrastructure as a code (IAC) solution is triggered to perform the self-healing action. Users, instead of guiding and facilitating this process, simply monitor its progress through data visualizations and dashboards, analyze the events in aggregate, and then design and build proactive improvements.

As appealing as this scenario sounds, building end-to-end scenarios one after the other is a risky approach. Investments in individual self-healing scenarios can become obsolete when the pattern no longer appears. For example, if a self-healing scenario of automatic reboots is built to address stuck threads due to coding issues, the investment in the scenario becomes obsolete once the code is fixed. A less risky, longer lasting result can be achieved by taking a multi-pronged approach that matures and optimizes the component capabilities on the path to lower the cost of building end-to-end scenarios. Taking this multi-pronged approach (Figure 1) will also help avoid common pitfalls in the AIOps journey such as high data management cost accumulation and delayed or short-lived

**Figure 1: AIOps: Putting the Pieces Together**

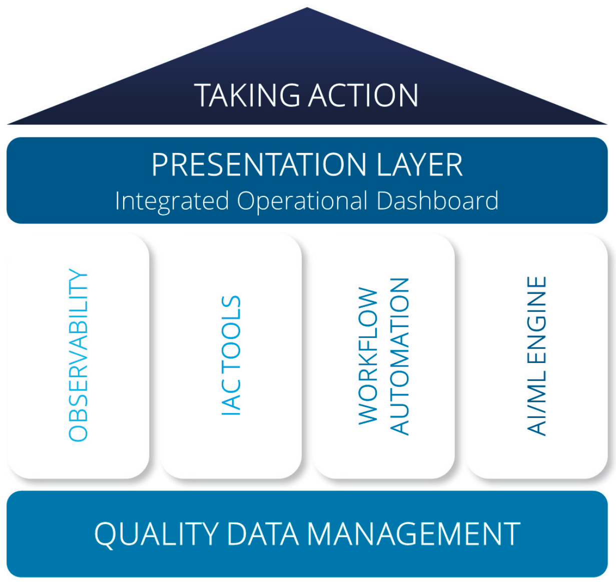




Maturing in AIOps – creating value at each step of the journey

# How to get started?

## Prioritize quality data management to limit ingest and storage costs

The "blocking and tackling" of any business intelligence (BI) project is data quality and data cleaning. The classic phrase of "garbage in and garbage out" is an homage to this notion. Data quality is an issue for AIOps, but it is more of a "noise in and noise out" than a "garbage in and garbage out" issue since AIOps data is mainly generated by machines that use standard formats. Further, storing noisy data incurs compute, storage, and licensing costs with little to no benefit.

Quality data management is a method of collecting data in increments, filtering as much as possible at the source and building smart technology to digest the data. This approach helps ensure that only the necessary data is collected for a specific use case or compliance issue. Quality data management involves using configured agents to filter at the source, using multiple storage tiers to control costs, and creating efficient indexes to search the data (Figure 2). AIOps depends on collecting and storing data that can be used for various purposes, including compliance, analysis, and decision-making. However, collecting and storing data can be an expensive and time-consuming process, especially if the data is not effectively managed. This is where quality data management comes into play.

For instance, when Deloitte began supporting a "4th estate" Defense client, we found that the monitoring process was decentralized, and data was held in separate vendor systems that were inaccessible to the newly created event management team.

**Figure 2: Quality Data Management**

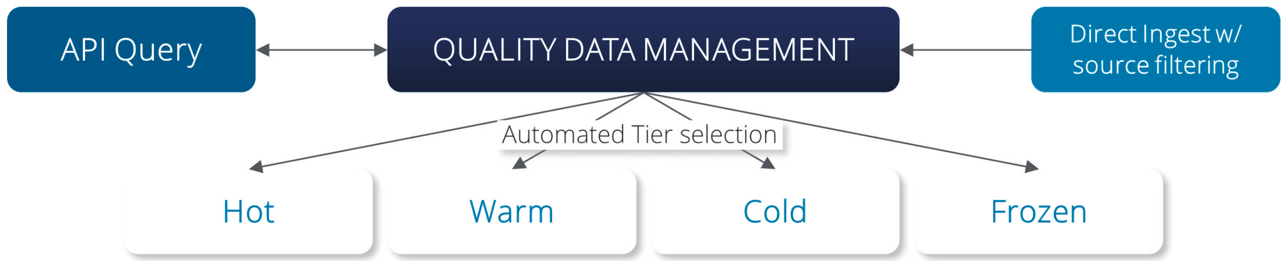


Initially, this team needed presentation layer capabilities to get a broad view of enterprise system health. Because centralization of event management was in its initial stages, investing in the collection and storage of duplicate data would have rapidly increased costs while the Deloitte team conducted a tool rationalization and selection initiative to determine the future mix of AIOps technologies. Alternatively, the Deloitte team responsible for tools queried data from the decentralized systems instead of collecting it. The team only staged data that was needed for the visualization use cases in the early sprints.

Further along in the journey as the observability platform and the AI / ML (Machine Learning) capabilities matured, the Deloitte team shifted to collect data into a centralized data lake for deeper analysis. As data accumulated, more sophisticated storage and retrieval techniques were needed to control the costs. Using tiered storage with hot, warm, cold, and frozen tiers moved data to more inexpensive disks and cloud storage-controlled costs. Similarly, developing an indexing strategy created more efficient retrieval for searches, dashboards, role-based access, and machine learning models also controlled costs. While collecting instead of querying data did increase costs overall, aligning the pace of these costs to the pace of the project's benefits maintained the position of value needed to sustain the journey.

# Build workflow automation for process efficiency

Both technical and administrative workflow automations increase output, improve quality, and lower costs. Technical automation involves the use of technology to automate processes such as infrastructure management, software deployment, and monitoring. This type of automation is typically handled by IT professionals and requires a deep understanding of the underlying technology.

Administrative automation, on the other hand, focuses on automating business processes such as data entry, invoice processing, and employee onboarding. This type of automation is typically owned by non-technical staff and requires a user-friendly interface that can be easily customized to meet the needs of each individual department.

This interim step can secure quick wins, reducing time and errors in routine tasks while server-side solutions are constructed

Ultimately, to create end-to-end scenarios, technical and administrative workflows need to converge. For example, forms to request compute or storage resources can be automated through a management process that accounts for spend and approves budget use. This process can then be tied to automated technical workflows that involve infrastructure as code solutions that will deliver the compute and storage resources requested through the administrative process. The execution of these processes will need to be monitored through observability tools and presented in a unified dashboard to confirm fulfillment.

**Figure 3: Automation Use Case Pipeline**

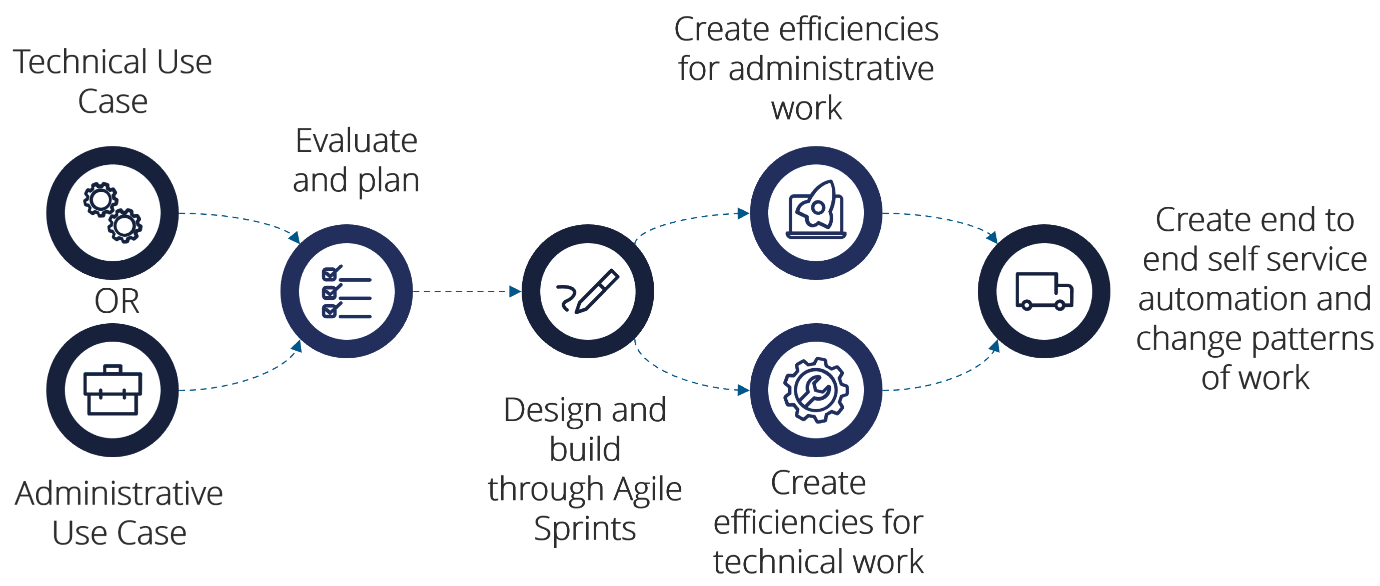




In both scenarios, the journey begins by clearly defining and standardizing the process steps before coding them into a solution. At these initial stages, benefits are realized from documenting and standardizing the workflows, which lay the foundation for later investments in coding and dashboarding.

An interim stage before full server-side implementation of technical or administrative workflow automation is to use a user-based desktop automation tool like those found in Blue Prism, UIpath, and others. These tools can act as bridge solutions while server-side solutions with complex authentication requirements are built.

WORKFLOW AUTOMATION

**Automated Ticket Management**

Deloitte delivered an automated Service Center Solution for Intelligence Community client which increased self-service metrics from 60% to 80% and decreased customer time to enter a ticket by 80%

Maturing in AIOps – creating value at each step of the journey

# Establish observability for faster recovery, preventative maintenance, and more resilient applications

Transparency into how applications function is the main purpose of any observability implementation. By gaining insight into the inner workings of applications, businesses can better understand how their systems operate and identify areas for improvement. However, "transparency" should not be the end goal. It is only when the data is used to affect architectures, transform the way assets are managed and reported, or lower the cost of compliance activities that its value is clear.

While data associated with observability software can validate or refute enterprise architecture documentation, and discovering the assets that exist in an enterprise can help with asset management, setting transparency as an end goal does not realize the underlying benefits of either of these scenarios. In fact, transparency as an end goal leads to the pitfall of high data management costs with delayed or short-lived benefits. Instead, the true value of observability lies in how the data is used to affect architectures and transform how assets are managed and reported. By leveraging the insights gained from observability, businesses can optimize their systems, improve performance, and enhance the overall user experience. This, in turn, can lead to lower costs, greater customer satisfaction, and stronger constituent services.

OBSERVABILITY

**Shifting Operations to the Left**

A federal civilian agency decided to embrace a "shift left" mentality, enabled by observability. They pushed operational and availability responsibility to product teams instead of a centralized incident management team as part of their move to the cloud. To support this strategic operating model, Deloitte constructed a centralized observability system to filter and route alerts relevant to the product teams.



Maturing in AIOps – creating value at each step of the journey

Initially, observability technology is used for up/down monitoring of applications and equipment, resource monitoring, and to visualize application telemetry data. Centralized monitoring through pings and synthetic tests are low-cost starting points to build the capabilities of event management that provide early warning indicators for customer-impacting incidents. Resource monitoring through standard protocols like SNMP and ping tests can help quickly organize IP address allocations and record when, say, a CPU is running out of capacity. Application performance monitoring using telemetry data to show how applications transact with infrastructure and other applications can inform enterprise architecture and reveal service issues. This type of metric collection supports early objectives to provide early warnings, prevent outages, and help with root cause analysis.

The next layer of maturity is where value begins to accelerate. Establishing more automation points and using AI and ML tools can help better understand the customer experience and pinpoint causes of outage and degradation through causal AI and time series analysis. These insights, however, need to be presented to application teams, architecture teams, and developers to spur projects focused on building more resilient systems and applications. Further, when converging with infrastructure as code and workflow automation technologies, automatic adjustments to architecture can be made to respond to degradation or even to marketplace changes. For example, if customer demand changes or cloud cost drivers change from one provider to the next, this combination of technologies can be used to automatically switch workloads to another cloud or a data center for service delivery. FinOps, advanced management of operational costs, cannot exist without data from observability systems, and can be made more agile through the convergence of AIOps technologies.

# Infuse infrastructure as code to administer, protect and stabilize at scale

While self-healing systems and automated provisioning are the high-water mark for AIOps, and of course, CIOs and IT leaders understand that you cannot do that without IAC tools, maturing these tools on the journey to self-healing can scale labor and increase availability. IAC tools, both cloud specific and general-purpose, are essential for the AIOps journey toward self-healing systems.

We all know that many outages at clients are caused by configuration or syntax errors, often known as "fat finger" errors, that too frequently occur during business hours. Early in the journey to AIOps, IAC tools can address these challenges by replicating "known to be good" test configurations at scheduled times. Combined with observability tools, the deployment of these configurations can be recorded automatically and displayed in dashboards and reports. In support of one-to-many configuration with automated deployments, IAC tools like Ansible allow for some key things. 1) The ability to collect known good lists of configuration items for efficiency and quality 2) The ability to detect configuration drift when, say, a developer makes a config change, that they shouldn't 3) The ability to make a configuration change on a mass scale for agility 4) A centralized repository to review configuration compliance (i.e., STIG and other mandates). These early use cases eliminate human error from common configuration changes and reduce the labor needed to deploy changes across the enterprise through one-to-many deployments.

Once an organization has benefited from one-to-many configuration deployments, the next step is to automate service provisioning. This involves integrating workflow automation into IAC tools that are also integrated with presentation layer tools. Dashboards show business and technical leaders their processes functioning, and workflows give them the opportunity to make decisions. However, before this can happen, administrative processes and governance must be standardized, and the organization must be designed to support them

with steering committees and cost accounting norms. Otherwise, integration, configuration, and coding will need to be redone as organizational processes and governance fluctuate. End-to-end automated service provisioning offers benefits such as proper tagging of resources, governance control gates through online systems, and fast deployments without systems administrator labor. This type of automation can accelerate FinOps programs that rely heavily on proper tagging of workloads and cost drivers. In addition, automated service provisioning sets the stage for advanced scenarios with IAC tools like self-healing, self- performance optimizing, and self-cost optimizing systems that require integration with an AI engine to enable decision making at volume with speed.

To achieve certain scenarios in AIOps, an AI engine is necessary. For instance, self-healing a system is only useful if the system can perform the tasks faster than a system administrator, or if it can act before a customer notices a problem. Performance optimization is beneficial only when it can handle a capacity issue promptly, and cost optimization requires an AI engine to present options to a FinOps manager to decide where to move workloads as costs fluctuate with cloud providers and customer demand. IAC tool maturity is necessary to adapt and respond to changing conditions in all these scenarios.

**IAC TOOLS**

**Detecting Configuration Drift**

While working with a Defense client, the Deloitte team integrated IAC tools with observability tools to detect drift from baseline STIG configurations. Scripts were created and run using open source SCAP tools to return reports to a centralized dashboard for audit reporting, and remediation actions.

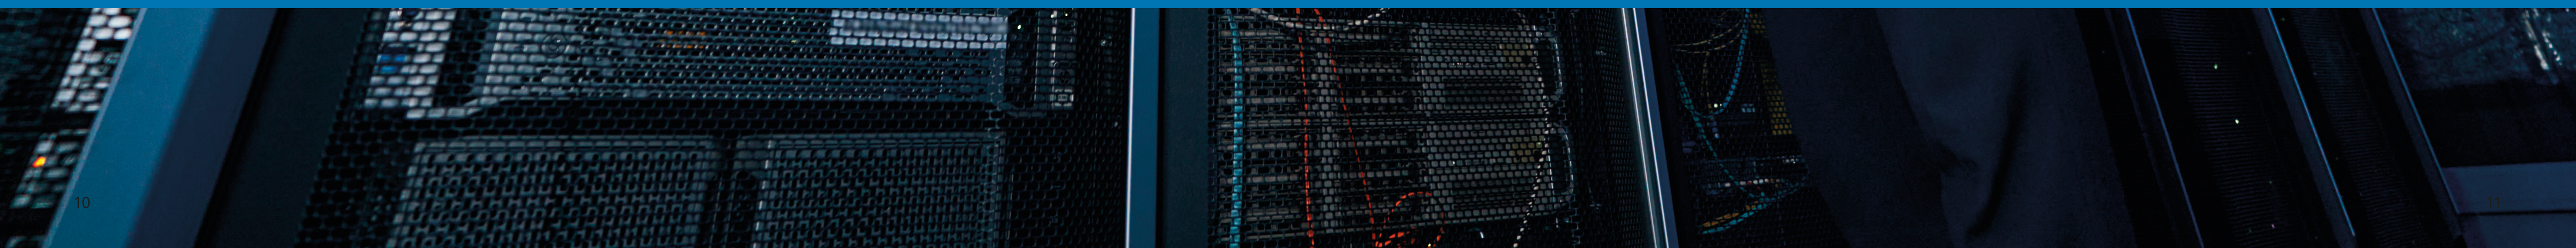

# Apply AI to perform the filtering and analysis that humans can't

As the complexity of IT systems continues to grow, the amount of data generated by monitoring and observability systems also increases exponentially. This data can include metrics like server performance, application response times, network throughput, and more. However, simply collecting this data is not enough - IT leaders need tools that can help them make sense of it all.

One approach is to use AI and ML algorithms to analyze this data for patterns, trends, and anomalies. This can help IT teams detect and diagnose issues more quickly, as well as identify opportunities for optimization and improvement. For example, an AI-powered observability tool could automatically generate a service map that shows how different applications and infrastructure components are interconnected. This can help IT teams understand the dependencies and relationships between different systems and identify potential points of failure or bottlenecks. Service maps and their graphical representations of interconnectedness that are coded and maintained solely by humans face a constant (and often losing) battle with obsolescence day-to-day as applications and infrastructures change.

However, it's important to note that AI can't replace human analysis entirely. While AI can be incredibly powerful at processing and analyzing large volumes of data, it still requires human oversight and input to help ensure the results are accurate and relevant. For instance, IT teams may need to manually tag different components of their infrastructure with metadata that describes their purpose, location, or security requirements. This information can then be used by AI algorithms to better understand the relationships between different systems and identify potential issues.

Overall, the combination of AI-powered observability tools and human expertise can help IT

leaders make smarter decisions and optimize their systems for peak performance. As the volume of data generated by IT systems continues to grow, these tools will become increasingly essential for managing and monitoring complex infrastructure.

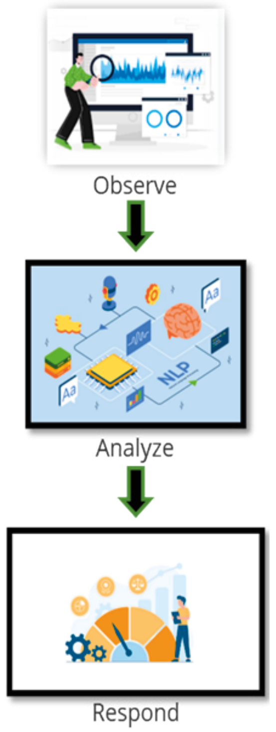


Early use cases of AI and ML apply them to make observability tools more relevant and useful. A common issue with monitoring and observability is noisy alerting, an issue in which too many alerts are received, and response teams are unable to distinguish between service impacting alerts, preventative maintenance alerts, and false negative alerts. Event and incident managers create checklists and plans of action for operational scenarios.



Maturing in AIOps – creating value at each step of the journey

Alert responses range from standing up a Priority 1 incident to creating a ticket to suppress a false negative. The AI tools integrated with most observability products can train models for "anomaly detection" and alert classification. Some models simply use historical data to point out observation outliers and other models can be trained with human feedback to classify alerts. These classified alerts can trigger workflows used for tickets, feedback, incidents, or provisioning.

Intermediate scenarios for AI and ML evolve correlation and causation capabilities. Leading observability tools create value by lowering mean time to recovery through service maps that show the technology used to provide business services. These maps use causal AI systems to pinpoint the component(s) responsible for an issue of outage or degradation. Known as root cause analysis, pinpointing the source of an issue whether in code or infrastructure is the first step in recovery and observability tools can make this an almost instantaneous diagnosis activity. The correlation happening in the background between time and component can be extended to data from workflow systems for change management or IAC systems to further understand the event that caused a failure.

As the field of artificial intelligence and machine learning continues to evolve, advanced scenarios that incorporate the latest developments in generative AI are becoming more practical. One of the key components is natural language processing (NLP) which has evolved into large language models. These technologies allow for the classification of language data and the identification of patterns and correlations. One area where this technology has already been applied is in customer service chat bots. By using NLP to analyze customer input, these bots can quickly identify the nature of the problem and suggest relevant knowledge articles that can help resolve it. This same approach can also be applied to machine alerts from observability systems, allowing for quick, efficient troubleshooting and problem resolution.

The potential applications of NLP and generative AI go far beyond simple correlations and knowledge article recommendations. With the right programming and data inputs, these systems can create new technical checklists, IAC playbooks, code changes, and even architectural changes.

This truly transformative capability has the potential to revolutionize not just IT operations, but many other fields as well. Of course, as with any powerful technology, there are also governance concerns to consider (e.g., architecture and sensitive data handling). Depending on the specific use case, it may be necessary to implement control gates or other oversight mechanisms to help ensure that the output of generative AI solutions is accurate, reliable, and aligned with organizational goals and ethics. However, with the right approach, these solutions can be a powerful tool for driving innovation, improving efficiency, and unlocking new insights and capabilities.

**AI / ML ENGINE**

**Smarter Alerts with AI / ML**

A state health exchange undergoing a multi-year modernization project uses AI and ML to reduce alert noise and boost collaboration between internal teams and external suppliers. The AI-powered insights also identified performance bottlenecks and helped optimize the end user experience.

# Create presentation layer dashboards to digest the journey

Getting better at all of this requires humans to understand and digest the monitoring data produced from workflows, IAC, observability, and AI. Charts, graphs, and visuals can provide the data and create widgets to drive insights for decision- making. This may sound like a use case for BI tools and in many ways, it is. BI tools as a product category, however, are biased to business or mission data and focus mostly on connecting to data sources leaving the transformation, normalization, and orchestration to the BI developer. Because much of the AIOps data ingested is standard and use cases are often very similar across industries, implementation at the presentation layer should be with a tool that has prebuilt plug-ins and a library of mashups and integrations out of the box.

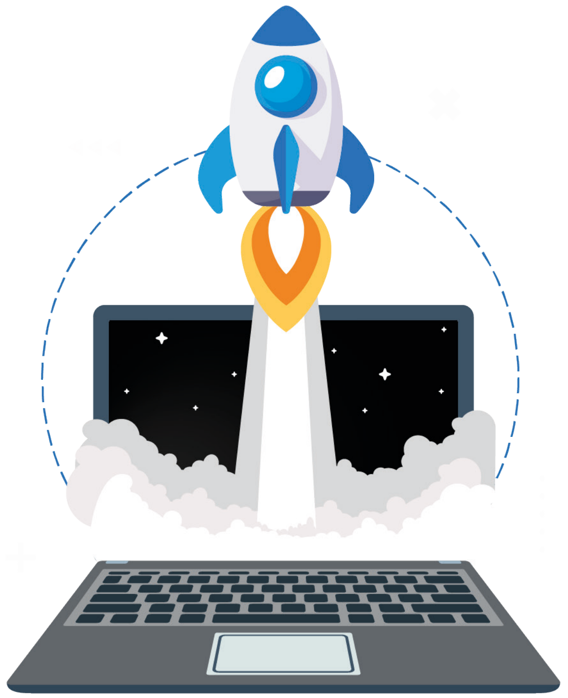

In early stages of maturity at the presentation layer, data will be decentralized in the vendor tools that ship with infrastructure products and in tools that are often included as part of the scope of custom software projects. During this initial stage, the central dashboarding tool primarily serves as a query tool to provide access to data and dashboards that are distributed across various sources. The focus is on categorizing and cataloging the available data sources to enable non-system administrators to access the data. At this stage, a reliable BI tool and/or JavaScript with an ETL tool are sufficient for these early dashboards.

However, as organizations progress beyond the initial proofs of concept and pilot dashboards, they need to make a sourcing decision. This decision considers tradeoffs between build vs. buy and picking a top-tier product vs. an emerging product from a cost modeling standpoint. From a capability standpoint, there are three categories of tools to consider.

The first category includes custom overlay tools that function on top of established platforms, such as Splunk or ServiceNow. The benefit of these tools are regular updates that are supported exclusively by the underlying platform.

The second category includes dedicated observability tools like Dynatrace that come bundled with visualizations. The easy-to-build visuals in these tools are geared towards technical team use cases and require effort to extend to other use cases.

The third category is tools designed to integrate data from multiple sources, combine it for visualizations and in some cases trigger action. These “managers of managers” that include PagerDuty, Grafana, and Deloitte’s own OperateEdge™, have the advantage of pre-built modules and visualizations for integrated scenarios, but can sometimes require custom code to meet every desired use case. To use any of these tools, data needs to be centralized. In this intermediate phase, organizations will still need to prioritize what data they centralize and when, to maintain quality data management.

In later stages, presentations can be extended to customers. It is a later stage since the components used to create simple-to-digest visualizations, like stoplight charts, must be tested and self- explanatory. False negatives and false positives can make customers quickly lose confidence in availability visualizations in particular.

**PRESENTATION LAYER**
**Integrated Operational Dashboard**

Multiple US State governments use Deloitte’s IT Operations Management platform, OperateEdge™ to integrate with observability tools, ITSM tools, DevSecOps tools, and to leverage AI to drive greater insight into all activities associated with effectively managing enterprise IT operations.

# Recommendations

If you ask a leading GenAI platform what the benefits of AIOps are, it concludes by saying that: "Overall, AIOps empowers organizations to streamline IT operations, enhance efficiency, and deliver better services to both internal and external stakeholders. It's important to note that while AIOps offers significant benefits, successful implementation requires careful planning, collaboration, and ongoing monitoring to achieve the desired outcomes." This underscores that AIOps is a journey that requires a guide.

The field of artificial intelligence has come a long way in recent years and the application of AI in operations has become increasingly popular. An effective AIOps guide can build the technology, reengineer the business process, create efficient operations and chart a course to go from simply performing operations to innovating the delivery of services. It is important to note that this journey requires patience and a long-term perspective. The implementation of AI in operations is not a quick fix solution, but rather a transformation that takes time and effort to yield results. While focusing on sophisticated scenarios may seem like the logical choice, it is crucial to also prioritize quick wins by maturing use of component capabilities. These smaller victories can provide momentum and support for continuing the journey towards this more impressive end to end innovations.

The framework for this journey can be found in Operate to Innovate, a Deloitte Methodology. Fundamentally it helps do the things that create better operations, tie them together, and make room for innovations like AIOps solutions. Also key to the journey is:

- **Building the right team**
- **Selecting the right mix of tools**
- **Maturing capabilities at a similar pace**
- **Managing the backlog and prioritizing the high value use cases**
- **Standardizing the existing processes**

The journey begins with an assessment of the current capabilities that can help build a roadmap for quick wins and end-to-end solutions. Just as governments manage program effectiveness, measuring the progress on the AIOps journey is needed to show the value and maintain the fuel in the journey to starship-like scenarios.

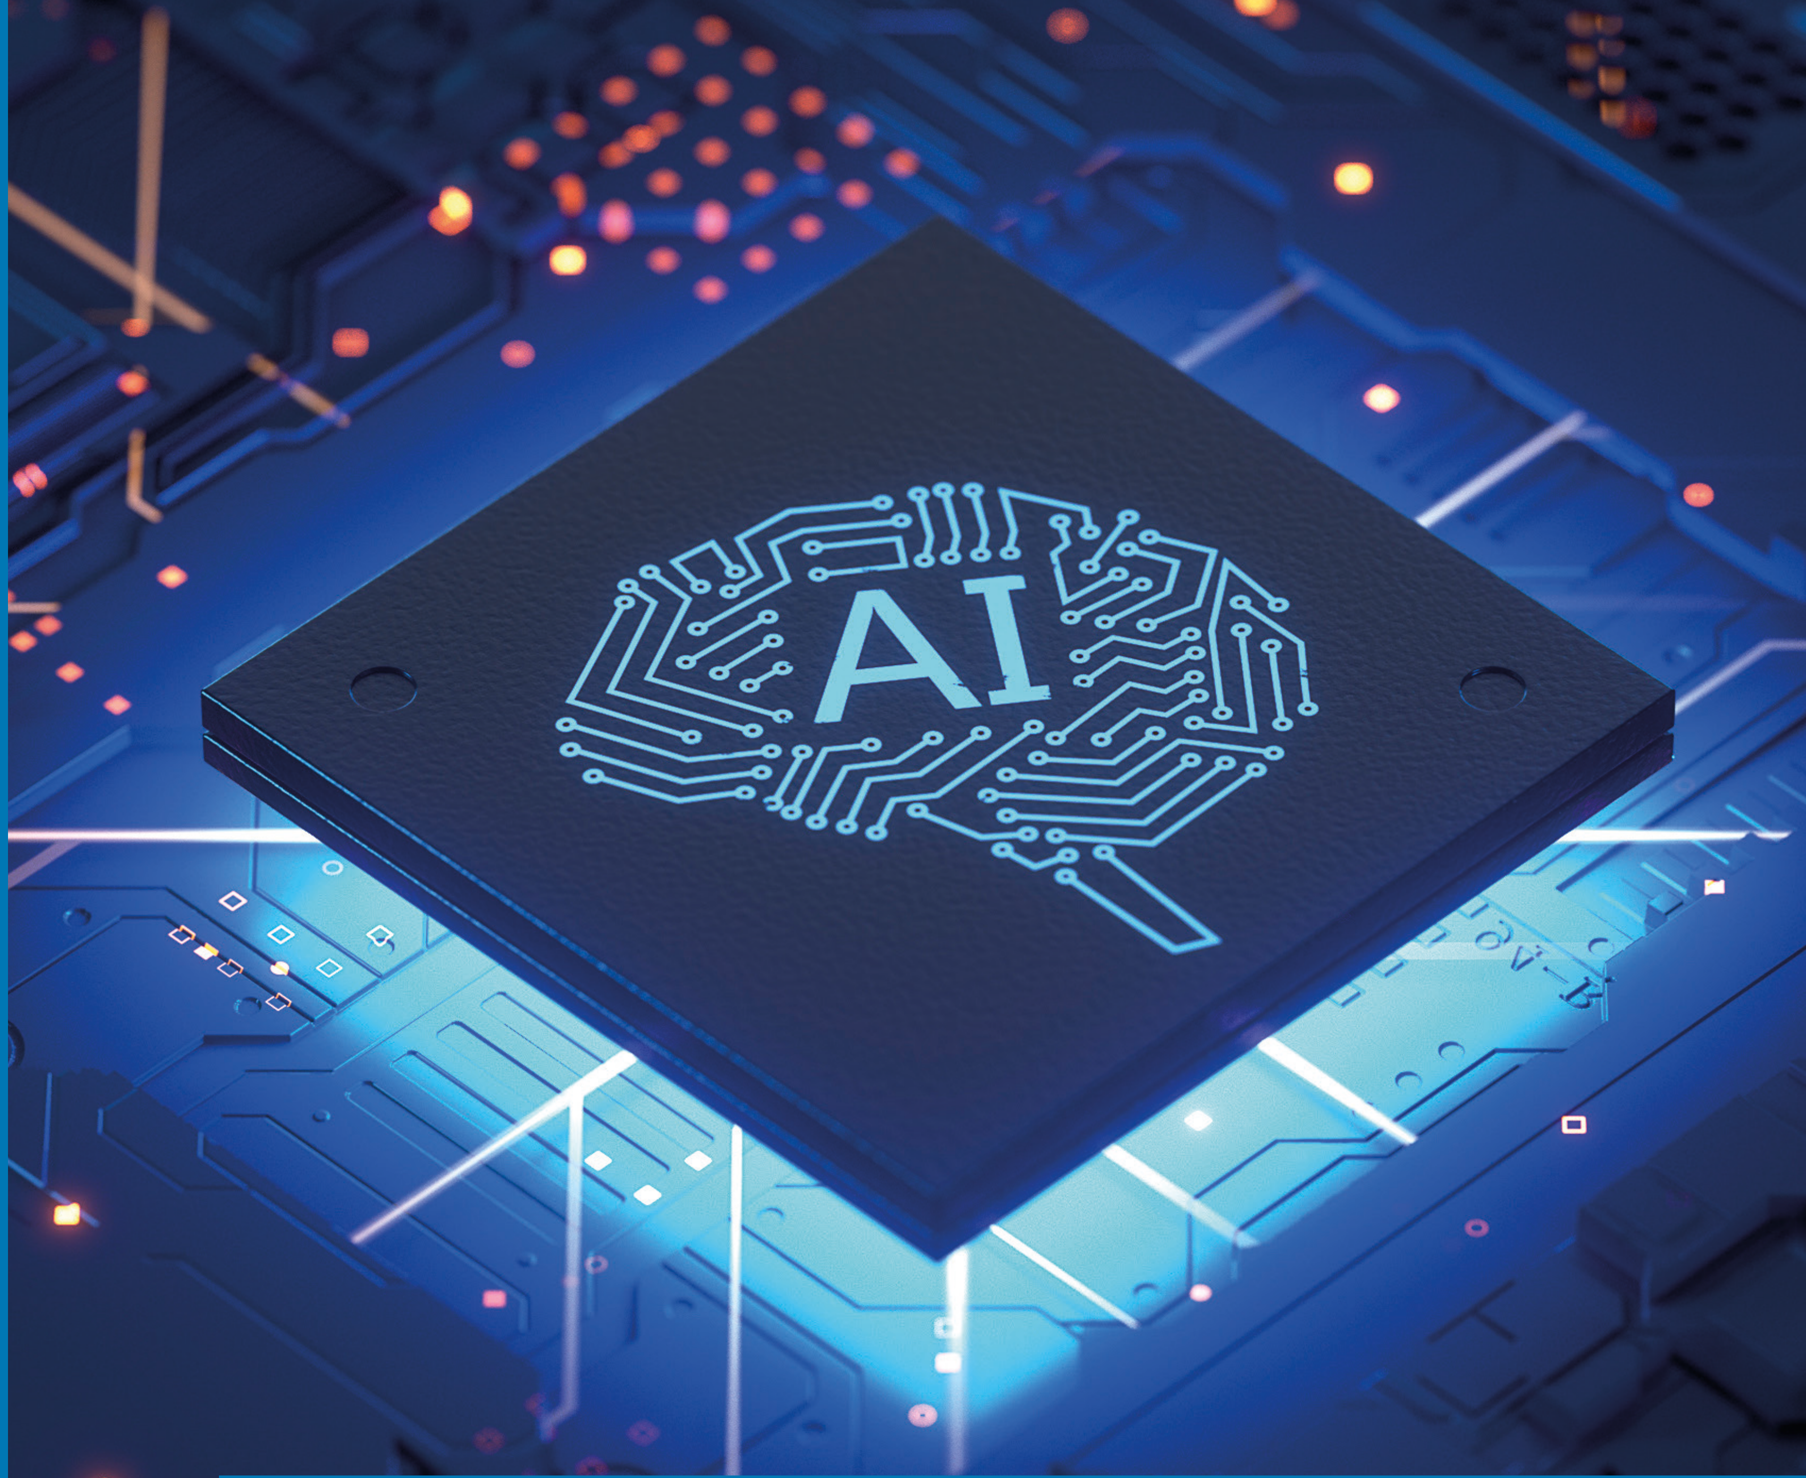


# Contributors


**Brian Pagliano**
Author
Core Technology Operations
bpagliano@deloitte.com

**Manny Cadre**
Author
Core Technology Operations
mcadre@deloitte.com

Designed by CoRe Creative Services. RITM1697164