

WHAT IS...?

AIOps

www.itsmacademy.com



WHAT'S INSIDE:

What is AIOps and What Are Its Origins?

Core Technologies

- ▶ Big Data
- ▶ AI and ML

AIOps and Operational Metrics

Implementing AIOps

It's a Journey!

Want to Learn More?

This course is derived with permission from the DevOps Institute's AIOps Foundation course. DevOps Institute is a register trademark of the PeopleCert group.

W What is AIOps and What Are Its Origins?

AIOps (Artificial Intelligence for IT Operations) represents a transformative approach to managing IT systems by leveraging advanced technologies such as artificial intelligence (AI) and machine learning (ML).

The term "AIOps" was coined by Gartner in 2016 and was first used to describe a rising industry category aimed at using machine learning to solve the problems arising from operating hyperscale cloud infrastructure. AIOps solutions leverage big data (BD) and ML capabilities to accelerate, automate, and simplify all tasks related to keeping the performance of systems running with high efficiency and reliability.

AIOps as a practice is very related to the preceding ITOA (IT Operations Analytics). A fundamental difference between the two is that ITOA leverages big data analytics alone, whereas AIOps is an evolution of those practices that uses ML *and* AI.

AIOps has emerged in response to three overlapping trends:

- **Increased infrastructure complexity and the cloud** as organizations migrate from monolithic applications to complex solutions built from numerous microservices and deployed to infinitely scalable environments
- **Data from increasingly observable systems** as development teams incorporate [observability](#) into their applications thus increasing data volumes exponentially
- **Broad access to mature machine learning capabilities** as ML and AI are rapidly becoming staple technologies, and ML capabilities are now present across most user experiences

ARTIFICIAL INTELLIGENCE

Software that thinks like a human through analysis and reasoning to perform complex tasks.

Continue reading to learn more...

The Underlying Business Driver

Digital transformation is the process of leveraging digital technologies to fundamentally change and improve business operations, customer experiences, and overall organizational performance. It involves adopting new technologies and embracing a digital-first mindset.

The adoption of digital technologies aims to deliver the following outcomes:

- Acceleration of innovation and invention
- Improvement of customer experiences
- The transformation of business processes

All these motives drive the need for new system architectures and the migration of services into the cloud. As these initiatives take place and systems are exposed to ever increasing scale, IT teams must adapt to maintain control of new, rapidly changing, massively scalable environments.

AIOps is not a silver bullet to any problem.

The outcomes organizations are looking to achieve must be realistic. AIOps can, however, provide benefits to a wide range of stakeholders including security, development, IT operations, site reliability engineering (SRE), and DevOps practitioners. Perhaps most importantly, it provides business benefits such as improved customer and employee experience, increased product and service reliability and availability, increased operational efficiency, decreased IT operations costs, real-time visibility into IT operations, the ability to make data-driven decisions, enhanced security, and enhanced risk mitigation and compliance management capabilities.

Core Technologies

AIOps solutions integrate several core technologies including big data and ML to enhance IT operations. These technologies make it possible to analyze vast amounts of data, learn from historical patterns, and generate actionable insights to support IT operations and decision-making.

These technologies integrate with other tools such as knowledge management and collaboration systems (e.g., chat) and leverage automation to free up humans to focus on more strategic and complex activities.

BIG DATA

Data sets that are too large or diverse to be dealt with using traditional data-processing techniques.

Big Data

For an AIOps solution, data is at the core. Data is produced by systems across the organization and throughout the software development lifecycle including:

- During build and release processes (e.g., CI/CD)
- During testing (e.g., security, performance, and integration testing)
- By applications on runtime across multiple environments
- By infrastructure, managed services, and monitoring and alerting systems
- From end users

Big data refers to large and complex data sets that are too voluminous, diverse, and fast-changing to be effectively managed and analyzed using traditional data processing techniques.

A big data dataset may be built from different types of data including:

- **Structured data** – data that is stored and processed in a fixed format with a clearly defined schema (e.g., dates and credit card numbers)
- **Semi-structured data** – data that contains tags, metadata, and formal semantics, but does not conform to a specific data model (e.g., tab-delimited files and log files)
- **Unstructured data** – data where the form and structure are not known (e.g., free form text or images)

Big data has a shelf life and gravity, and it comes from many different sources of varying types. The value of the data being produced (i.e., its **shelf life**) decreases in value over time. Extracting data value quickly is a key objective of systems fed by big data. **Gravity** refers to the fact that large data sets attract other data, applications and use cases.

The **5 Vs** of big data are a framework that describes the key characteristics and challenges associated with big data. They are:

- **Volume** – the vast amount of data that is generated and collected.
- **Velocity** – the speed at which data is generated, processed, and analyzed.
- **Variety** – the diverse types and formats of data that are encountered in big data environments.
- **Veracity** – the quality, accuracy, and reliability of data.
- **Value** – the ability to extract meaningful insights and value from big data.

In AIOps, data is the backbone that enables the system to function effectively. An AIOps solution leverages diverse data sources to provide comprehensive monitoring and analytics. Essential data types (known as **MELT**) typically include:

- **Metrics** – quantitative data that provides measurements over time, such as CPU usage, memory consumption, and network throughput.
- **Events** – discrete occurrences within the IT environment, such as user logins, error messages, and configuration changes.
- **Logs** – detailed records of activities and transactions, often used for troubleshooting and identifying patterns.
- **Traces** – data that follows the path of transactions across different systems and components, useful for pinpointing performance bottlenecks and understanding system interactions.

From Source to AIOps

The journey from raw data to actionable insights in AIOps involves several key steps, facilitated by data pipelines. A data pipeline oversees taking data from its source, handling the required tasks to make the data clean and reliable for use by ML, and putting it into a target data repository that the AIOps solution can use.

These pipelines perform crucial tasks including:

- **Extraction and collection** – getting the data from its source in close to real time.
- **Enrichment and filtering** – filtering any data that is not relevant and enriching relevant data with additional context for its value to increase.
- **Cleaning and integration** – cleaning up, deduplicating, and integrating data into a comprehensive data model.
- **Storage** – storing the data in a repository that can be efficiently accessed by AIOps ML processes.

The ultimate goal of analyzing big data is to derive actionable insights, make informed decisions, and create business value. The challenge lies in identifying and extracting valuable insights from these massive and diverse data sets. Here is where technologies such as AI and ML are used to uncover patterns, trends, and correlations within the data.

Artificial Intelligence and Machine Learning

The terms 'artificial intelligence' and 'machine learning' are at times used interchangeably but they are not the same thing.

Artificial intelligence

AI is a broad concept that encapsulates various technologies and techniques. At a high level, AI can be described as: "The replication or simulation of human intelligence by computer systems."

There are generally three types or levels of AI:

- **Narrow AI (also known as "weak AI")** is task-specific AI such as facial recognition with limited parameters.
- **General AI (also known as "strong AI")** is AI that can solve unplanned, complex problems.
- **Super AI** is AI that surpasses human abilities and performance.

To date, most applications of AI encountered in daily life fall under the category of narrow AI.

Machine Learning

ML is a subset of AI that uses algorithms to recognize patterns from large amounts of data. It uses massive amounts of historical data to generate accurate inferences or predictions.

MACHINE LEARNING

Data analysis that uses algorithms to recognize patterns from large amounts of data to learn.

ML models imitate learning from data, and gradually improve their accuracy through iterative training and optimization processes. An **ML Model** is a program that can find patterns or make decisions from a previously unseen dataset.

ML uses data to train models, just as humans learn based on past experiences. It is based on this data that the model will be able to "learn" how to interpret and infer outputs, depending on future inputs.

There are two approaches to providing data for training a model: supervised and unsupervised.

- **Supervised learning** requires a higher degree of human involvement, since the model will require sample data that has been labeled for input as well as for output.
- **Unsupervised learning** works with unlabeled data and requires a lower degree of human intervention.

Unsupervised learning models are given raw data as input without explicit human input, and the machine works to autonomously identify patterns and structures from that raw data.

AIOps applies AI to IT operations. It uses ML to make operating IT systems more efficient and reliable and require less human effort.

Just as AI and ML are not the same, it is important to understand the distinction between AIOps and MLOps. MLOps applies patterns introduced by DevOps towards operationalizing ML models. It implements CI/CD, testing, and accelerated development lifecycles to the ML model development.

MLOps can be confused with ModelOps, which is a set of capabilities that primarily focuses on the governance and the full life cycle management of all AI and decision models.

Simply put:

- AIOps is AI for IT operations
- MLOps focuses only on the operationalization of ML models
- ModelOps focuses on operationalizing all AI and decision models

AIOps and the Future of AI

Generative AI is revolutionizing how AI can partner with humans, venturing into the territory of creative problem solving. **Generative AI** is a type of ML model that uses big data sets to learn how to create (generate) new data. Large language models (LLMs) are core to the capabilities of generative AI. These models are built using extremely large data sets and are designed to understand, generate, and manipulate human language.

In the context of AIOps, generative AI can provide significant enhancements by automating and optimizing various aspects of IT operations. Examples (provided by ChatGPT) include:

- When an issue is detected, AI can generate and execute scripts or actions to remediate the problem without human intervention.
- By analyzing log files, performance metrics, and past incident reports, AI can predict when a server is likely to fail and generate a maintenance schedule or alert the IT team to take preventive measures.
- After resolving an incident, AI can automatically generate a detailed incident report and update the knowledge base with steps taken for resolution, potential causes, and preventive measures.

Generative AI can even help users learn how to use AIOps software. thus reducing the learning curve and increasing user adoption.

AIOps and Operational Metrics

According to Gartner, one of the main barriers to implementing AIOps solutions is the difficulty measuring their value and understanding their benefits.

Understanding and tracking the performance of an AIOps solution is critical for the solution to:

- Produce the desired outcomes
- Simplify operations
- Reduce costs and eliminate waste
- Continuously improve its value to the organization

Tracking the performance of an AIOps solution involves monitoring various key metrics and indicators to assess its effectiveness in supporting IT operations.

For example, AIOps should result in improvements to a system's **RAMS** which stands for:

- **Reliability** – how long a resource can perform as expected without interruption
- **Availability** – portion of time a resource is ready to perform a requested action
- **Maintainability** – how quickly and easily something can be fixed
- **Safety** – a system's ability to not create harm throughout its lifecycle

AIOps can have a significant impact on various incident-related metrics, leading to improvements in incident response and resolution. Metrics related specifically to incidents include:

- **Mean time to resolve (MTTR)** – time it takes to recover a system to its fully operational state
- **Mean time between failures (MTBF)** – timespan that occurs between incidents in a system (often used to measure reliability)
- **Mean time to detect (MTTD)** – time it takes to detect an issue
- **Mean time to acknowledge (MTTA)** – time that passes between the first alert of an incident in the system and the first attempt to recover from this incident

Successful AIOps implementations can support DevOps teams in improving widely used DORA metrics over time including (as defined in the Accelerate State of DevOps Report):

- **Change failure rate** – how frequently a deployment introduces a failure that requires immediate intervention
- **Change lead time** – how long it takes a change to go from committed to deployed



AIOps can produce a direct benefit towards the objectives of SRE initiatives through outcomes such as:

- Eliminating toil
- Proactive SLA and SLO tracking
- Enhanced error budget definition and burn rates

AIOps can also simplify monitoring for compliance with internal policies and applicable laws.

As these metrics illustrate, different teams within an organization such as security, development, operations, SRE, and DevOps, can all derive benefits from AIOps. Clearly articulating the tangible benefits and potential outcomes of implementing AIOps to these stakeholders at all levels of the organization is essential.

Implementing AIOps

Embarking on an AIOps journey requires a clear strategy, with clearly defined, realistic desired outcomes. Lacking such a strategy is one of the most common causes for AIOps initiatives to fail. Documenting and distributing the strategy will help to gain buy-in and acceptance from all involved stakeholders.

It is critical to define clear objectives that will be used to validate the success of the AIOps implementation or identify improvement opportunities.

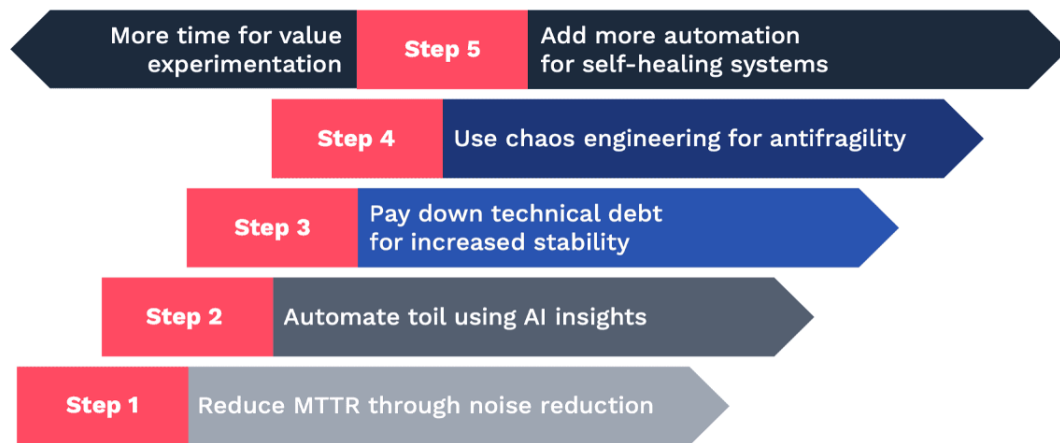
Important questions to consider include:

- Which indicators are you looking to improve?
- Which systems should be positively impacted by the implementation?
- Which teams should see the benefit of AIOps over time?
- What is the planned timeline to see results?

All these objectives should be documented and agreed upon by all our stakeholders and clearly understood by the implementation team. It is key is to ensure that expectations are clearly set and are aimed at a gradual improvement of quantifiable outcomes.

AIOps Capability Scale

The AIOps capability scale enables organizations to assess their current AIOps capabilities, identify gaps, and define a roadmap for advancing their AIOps implementation. This scale explains how teams gain increasing benefits as they evolve their use of AIOps technologies.



Initially, organizations rely on traditional monitoring tools and manual processes to identify and resolve issues. This approach typically results in high volumes of unplanned work, burnout, and alert fatigue. The implementation of AIOps can help reduce alert fatigue through noise reduction by deduplicating and correlating alerts and events. This greatly reduces both the MTTD and MTTR of incidents, along with the amount of time spent on unplanned work. That reclaimed time can be allocated to activities such as reducing toil and paying down technical debt.

As teams pay down technical debt, systems become more stable and the teams gain the confidence to plan controlled experiments that make it possible to understand how systems behave in the face of failure (i.e., chaos engineering). Teams can now devote even more time to automating fixes and achieving self-healing systems. Ultimately, this leads to teams having more time to focus on value experimentation.

To achieve even the initial outcomes, data quality is a crucial factor. Organizations shouldn't dedicate time to implementing machine learning until solid data is available. Finding the right balance of system coverage is also crucial. While organizations may want to start with a relatively small scope and set of objectives, the scope cannot be so small that it would not provide enough data points to allow ML to be of benefit.

AIOps and Organizational Culture

Creating a culture that supports AIOps adoption is crucial for successful implementation and long-term benefits. Adopting AIOps and the changes of mindset it introduces can be successfully handled by organizations that have already adopted other modern paradigms successfully.

For example, the culture of an organization must be well along the DevOps path, and there must be a good comfort level with distributed systems, cloud native systems design, and other paradigm shifts (e.g., moving from passive monitoring into observability) to help with the implementation of AIOps.

By nurturing a culture that embraces AIOps, organizations can create an environment where teams are empowered to leverage its capabilities effectively. This cultural shift promotes collaboration, innovation, and data-driven decision-making, leading to improved IT operations, efficiency, and business outcomes.

The Path to Implementation

For an AIOps implementation to succeed, the team must have access to subject matter experts across several different domains including:

- Infrastructure and cloud operations
- Site reliability engineering
- DevOps
- Data engineering
- Machine learning

Given the broad impact of a long-term AIOps strategy, it is important to ensure executive sponsorship and alignment exists.

Activities along the path to implementing AIOps include:

- Identifying clear, actionable objectives
- Building the right team
- Ensuring metrics exist to quantify progress
- Identifying data sources and ML models
- Architecting the data and AIOps engine
- Deploying the first iteration
- Measure and improve

As objectives are reached and the system matures, organizations can increase the scope by introducing new systems and data sources, and by extending the range of desired benefits to be derived from the initiative.

It's a Journey!

AIOps like any technology-driven approach to IT management is only as valuable as the outcomes that it helps to achieve. An organization's culture, how AIOps solutions are implemented, and how its benefits are communicated will all affect its perceived value.

It is necessary at all times to remember the objectives to be accomplished and *why* they need to be accomplished. A clear strategy with clearly defined and realistic desired outcomes is essential as are taking an iterative approach and having a commitment to continual improvement.

AIOps is a journey. It may take time to fully implement and realize its benefits. Regularly evaluate and adjust the strategy based on feedback, evolving technologies, and changing business needs to ensure ongoing success.

Want to Learn More?

Training helps individuals and organizations build and maintain their capabilities. Training also provides individuals the knowledge, skills and information needed to fill their role in the organization or achieve their career goals, along with a place to test and develop the confidence to use these skills in the workplace.

The **AIOps Foundation** course introduces the origins of AIOps and its key terms, basic concepts, and core technologies, along with strategies for implementing AIOps.

See ITSM Academy [class options and dates](#) or download the [course catalog](#).

Additional Resources:

- [ITSM Professor Blog](#) - a WEALTH of knowledge published weekly since 2008
- [Webinar Archives](#) - Monthly since 2007
- [ITSM Academy Resource Center](#)



ITSM Academy

We are a female owned small business, established in 2004. Our extensive catalog contains accredited and sustainable IT Service Management (ITSM) education and advice including; ITIL®, DevOps, Process Design (CPDE), Agile, Site Reliability Engineering (SRE), Value Stream Mapping (VSM) and Experience Level Agreement (XLA). Our business values are founded on trust, loyalty, professionalism and long term relationships.

...educate and inspire is not just our corporate slogan, it speaks to our core mission and goal.



Follow our founder and CXO, Lisa Schwartz, on [LinkedIn](#).

Instructors

Every ITSM Academy instructor is certified to the highest levels in the areas they train. They have years of hands-on IT practitioner experience, enabling them to effectively intertwine theory and real-life stories and scenarios. Using the highest quality content, this engaging training style encourages active group participation, allowing all learners to bring from class a wealth of practical and actionable knowledge.

Accreditations

All of ITSM Academy's certification courseware is developed or enhanced in-house and is accredited by independent, international organizations where applicable.



Game On! - Interactive Learning

Involves students in active learning, using the engaging qualities of a game, fueled by our subject matter experts.

Courseware Licensing (all developed or enhanced in house)

In addition to our public and corporate/onsite training, our courseware is available for licensing / co-branding under our flexible licensing program, including Train-the-Trainer (for qualifying organizations).

my.itsmacademy.com (digital portal)

Extends the learning experience with games, videos, exercises, sample exams, and course materials. It also provides instructors a vast repository of information and guidance to successfully prepare for and teach our courses.

Professional Education Hours (CPDs/PDUs/CPEs/CEUs):

ITSM Academy is proud to make it possible for individuals who attend our classes to earn professional education hours. (e.g., CPDs, PDUs, CPEs, CEUs). These professional education hours can be submitted to associations such as PeopleCert, the Project Management Institute and ISACA, if applicable.



[Read More](#)

The Story of the Academy

Today, ITSM Academy is widely recognized for its expertise in multiple IT frameworks (ITSM, ITIL, Process Engineering (CPDE), DevOps, Agile Service Management, Lean) and, more importantly, how they work together. But that's not where we started.

+1-954-491-3442

<http://itsmacademy.com>