



Hands-on AIOps

Best Practices Guide to Implementing
AIOps

—

Navin Sabharwal
Gaurav Bhardwaj

Apress®

Hands-on AIOps

Best Practices Guide to Implementing AIOps

Navin Sabharwal
Gaurav Bhardwaj

Apress®

Hands-on AIOps: Best Practices Guide to Implementing AIOps

Navin Sabharwal
New Delhi, India

Gaurav Bhardwaj
New Delhi, India

ISBN-13 (pbk): 978-1-4842-8266-3

<https://doi.org/10.1007/978-1-4842-8267-0>

ISBN-13 (electronic): 978-1-4842-8267-0

Copyright © 2022 by Navin Sabharwal and Gaurav Bhardwaj

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr

Acquisitions Editor: Celestin Suresh John

Development Editor: Laura Berendson

Coordinating Editor: Shrikant Vishwakarma

Copyeditor: Kim Wimpsett

Cover designed by eStudioCalamar

Cover image by Pixabay (pixabay.com)

Distributed to the book trade worldwide by Apress Media, LLC, 1 New York Plaza, New York, NY 10004, U.S.A. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail booktranslations@springernature.com; for reprint, paperback, or audio rights, please e-mail bookpermissions@springernature.com.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at www.apress.com/bulk-sales.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub (<https://github.com/Apress>). For more detailed information, please visit www.apress.com/source-code.

Printed on acid-free paper

*Dedicated to family, friends, mentors, and
Almighty God who has been there right from the beginning,
giving strength and skill.*

Table of Contents

About the Authors.....xi

About the Technical Reviewerxiii

Acknowledgments xv

Preface xvii

Introduction xxiii

Chapter 1: What Is AIOps? 1

 Introduction to AIOps..... 1

 Data Ingestion Layer..... 6

 Data Processing Layer..... 7

 Data Representation Layer 7

 AIOps Benefits..... 13

 Summary..... 17

Chapter 2: AIOps Architecture and Methodology 19

 AIOps Architecture 19

 The Core Platform 20

 Big Data 20

 Machine Learning..... 22

 The Three Key Areas in AIOps 23

 Observe 24

 Engage..... 41

 Act 49

TABLE OF CONTENTS

Application Discovery and Insights58

Making Connections: The Value of Data Correlation59

Summary.....60

Chapter 3: AIOps Challenges61

Organizational Change Management.....62

Monitoring Coverage and Data Availability63

Rigid Processes64

Lack of Understanding of Machine Learning and AIOps65

Expectations Mismatch.....66

Fragmented Functions and the CMDB67

Challenges in Machine Learning.....68

Data Drift.....69

Predictive Analytics Challenges70

Cost Savings Expectations.....70

Lack of Domain Inputs71

Summary.....72

Chapter 4: AIOps Supporting SRE and DevOps73

Overview of SRE and DevOps.....74

SRE Principles and AIOps.....78

 Principle 1: Embracing Risk.....79

 Principle 2: Service Level Objectives.....79

 Principle 3: Eliminating Toil.....80

 Principle 4: Monitoring81

 Principle 5: Automation.....83

 Principle 6: Release Engineering.....84

 Principle 7: Simplicity.....86

AIOps Enabling Visibility in SRE and DevOps	86
Culture	86
Automation of Processes	87
Measurement of Key Performance Indicators (KPIs)	87
Sharing.....	88
Summary.....	88
Chapter 5: Fundamentals of Machine Learning and AI.....	89
What Is Artificial Intelligence and Machine Learning?	89
Why Machine Learning Is Important	91
Types of Machine Learning	93
Machine Learning.....	93
Supervised (Inductive) Learning.....	94
Unsupervised Learning.....	96
Reinforcement Learning	99
Differences Between Supervised and Unsupervised Learning	100
Choosing the Machine Learning Approach	101
Natural Language Processing	102
What Is Natural Language Processing?	103
Syntactic Analysis.....	104
Semantic Analysis	105
NLP AIOps Use Cases	106
Sentiment Analysis	106
Language Translation	107
Text Extraction	107
Topic Classification.....	107
Deep Learning.....	108
Summary.....	112

TABLE OF CONTENTS

Chapter 6: AIOps Use Case: Deduplication113

Environment Setup..... 113

Software Installation 114

Launch Application..... 115

Performance Analysis of Models..... 116

 Mean Square Error/Root Mean Square Error 117

 Mean Absolute Error 119

 Mean Absolute Percentage Error 119

 Root Mean Squared Log Error 120

 Coefficient of Determination-R2 Score..... 120

Deduplication 121

Summary..... 135

Chapter 7: AIOps Use Case: Automated Baselineing137

Automated Baselineing Overview 137

Regression 139

 Linear Regression..... 140

Time-Series Models 149

 Time-Series Data..... 150

 Stationary Time Series 152

 Lag Variable 152

 ACF and PACF 153

 ARIMA..... 156

 Model Development..... 157

 SARIMA..... 161

 Implementation of ARIMA and SARIMA 162

Automated Baseline in APM and SecOps	185
Challenges with Dynamic Thresholding	185
Summary.....	186
Chapter 8: AIOps Use Case: Anomaly Detection	187
Anomaly Detection Overview	187
K-Means Algorithms.....	188
Correlation and Association	207
Topology-Based Correlation	208
Network Topology Correlation.....	209
Application Topology Correlation	211
Summary.....	215
Chapter 9: Setting Up AIOps	217
Step 1: Write an AIOps Charter.....	218
Step 2: Build Your AIOps Team	218
Step 3: Define Your AIOps Landscape	219
Step 4: Define Integrations and Data Sources	221
Step 5: Install and Configure the AIOps Engine	222
Step 6: Configure AIOps Features	224
Step 7: Deploy the Service Management Features	225
Step 8: Deploy Automation Features.....	226
Step 9: Measure Success.....	226
Step 10: Celebrate and Share Success	228
Guidelines on Implementing AIOps	229
Hype vs. Clarity.....	229
Be Goal and KPI Driven.....	229
Expectations	229
Time to Realize Benefits.....	230

TABLE OF CONTENTS

One Size Doesn't Fit All230

Organizational Change Management230

Plan Big, Start Small, and Iterate Fast.....231

Continually Improve.....231

The Future of AIOps.....231

Summary.....235

Index.....237

About the Authors



Navin Sabharwal has more than 23 years of industry experience. He is a thought leader and head of strategy and development in the areas of AgileOps, cloud computing, DevSecOps, AIOps, FinOps, artificial intelligence, and SaaS product engineering. He holds seven patents in the areas of AI and machine learning. He has authored 15+ books in these domains in the areas of cloud capacity management, observability, DevOps, cloud automation, containerization, Google AutoML, BERT, and NLP. He can be reached at Navinsabharwal@gmail.com and www.linkedin.com/in/navinsabharwal.



Gaurav Bhardwaj is a seasoned IT professional and technology evangelist with more than 17 years of expertise in service assurance, enterprise monitoring, event management, cloud computing, AI/ML-based software product development, engineering, and data analytics. He has international experience in solution designing and execution of IT automation strategies that are aligned with business goals. Gaurav has a proven track record of achievements in the roles of enterprise architect and consultant on IT transformation projects

ABOUT THE AUTHORS

for large and complex global engagements (including multiple Fortune 500 companies), migrating from legacy processes and platforms to next-generation IT environment powered by cloud-native and containerized apps, SDI and AIOps, and DevOps methodologies. He can be reached at Gaurav.Bhardwaj22@hotmail.com and www.linkedin.com/in/gauravbhardwaj-1.

About the Technical Reviewer



Pradeepta Mishra is the head of AI at FOSFOR by L&T Infotech (LTI), leading a group of data scientists, computational linguistics experts, and machine learning and deep learning experts in building artificial intelligence products. He was awarded “India’s Top 40Under40DataScientists” by *Analytics India Magazine* for two years in a row, 2019 and 2020. As an inventor, he has filed five patents

in different global locations. He is the author of five books, and his first book has been recommended by the HSLS Center at the University of Pittsburgh, Pennsylvania. His fourth book, *Pytorch Recipes*, was published by Apress and added to the Buswell Library in Illinois. His fifth book, *Practical Explainable AI Using Python*, was recently published by Apress and has been recognized as a textbook by the Barcelona Technology School (BTS) for its big data analytics course. He delivered a keynote session at the Global Data Science conference in California in 2018 and has delivered 500+ tech talks on data science, ML, DL, NLP, and AI in various universities, meetups, technical institutions, and community-arranged forums. He is a visiting faculty and academic advisory board for AI/ML at Reva University in Bangalore, India, as well as at various other universities. He has mentored and trained more than 2,000 data scientists and AI engineers in the past decade.

Acknowledgments

To my family, Shweta and Soumil, for being always there by my side and letting me sacrifice their time for my intellectual and spiritual pursuit. This and other accomplishments of my life wouldn't have been possible without your love and support.

To my mom and my sister for your love and support as always: without your blessings, nothing is possible.

To my coauthor Gaurav and the team at HCL who have been a source of inspiration.

Thank you to Celestine and the entire team at Apress for turning our ideas into reality. It has been an amazing experience authoring with you, and over the years, the speed of decision-making and the editorial support has been excellent.

—Navin

To my mom and dad for always being there by my side, giving invaluable support and encouragement and continuously striving for knowledge and wisdom.

To my wife and lovely daughters, Vaidehi and Maanvi, for being inspirations to me and for sacrificing their time.

To my coauthor and mentor, Navin, thank you for your guidance, support, and feedback.

To my teammates, Vipul Tyagi and Rajesh Upadhyay, thank you for your technical support.

Thank you, Lord, for giving me wisdom and strength through every aspect of life.

—Gaurav

Preface

Before starting our AIOps journey, let's briefly discuss automation and how it has evolved over the last decade.

Automation in the technology domain is defined as a system where a process or task can be performed with minimal human supervision and action. Humans have been automating tasks forever, and eventually mechanical machines that reduced human effort and increased efficiency were invented to reduce human effort. Today things such as manufacturing, which were partly automated earlier, are moving to 3D printing, which completely automates the process of manufacturing; however, designing what to manufacture is still in the human domain.

Thus, automation reduces human effort and uses machines or software to complete definable, repeatable tasks.

In the IT domain, there are various tasks that humans perform, including envisioning a new product or application, developing the software that translates these requirements into working software, and deploying infrastructure and applications and keeping them updated through their lifecycle.

IT teams have used automation extensively in every area, from software development to operations; however, this has largely been siloed and done without a formal system or method to automation. IT teams have used scripts, runbook automation tools, job scheduling systems, and robotic process automation systems to automate their tasks. These tools have resulted in increased efficiency and reduced the human requirement to operate IT environments.

With the increased adoption of cloud computing and DevOps principles, the provisioning of infrastructure-as-a-service and platform-as-a-service environments has also been automated, as has the deployment

PREFACE

of applications. This has resulted not just in automating the tasks and increasing efficiency but also in agility and speed, which provides businesses with the support to pivot and adapt to changing market needs by quickly changing the functionality and features based on customer and market feedback.

IT runs on three pillars: process, people, and technology. To be able to automate, one needs to be aware of the inter-relationships between these pillars. People use defined processes to work on technology, and with automation we are essentially automating the current processes that people are using to operate an environment. However, with the changing technology landscape and increased adoption and maturity of AI and machine learning capabilities, we can now look at the current processes and formulate new ones to leverage the transformational capabilities provided by these technologies. The current processes were set up with state-of-the-art technology at that point in time, and these processes then defined how humans should operate within that process to execute tasks to accomplish a goal; however, with a drastic change in the technology landscape, the processes need to change and adopt. As an example, with the cloud becoming prevalent, the IT processes need to change and adapt, and the sequential, nonautomated processes need to change to cater to the new capabilities such as autoprovisioning. In 2013, I talked about how capacity management would drastically change in the cloud computing world and that new procedures for cloud cost management would be required. Some of these concepts have been expanded to cover the entire financial operations piece under the umbrella of *cloud FinOps*.

Similarly, IT operations automation had existed in siloes for all this while. People used scripts, monitoring tools, runbook automation, configuration and deployment automation, and RPA tools and automated service management processes using ITSM tools. However, all that was happening under different domains, all getting integrated with siloed integrations.

On the technology front, AI and machine learning technologies became mainstream and were being used heavily in all aspects of business-facing and customer-facing applications from websites to search engines to collaboration and communications tools. AI took over the world of IT quickly; however, IT was late to adopt these technologies. While IT teams were using these technologies to create new applications with AI capabilities to customers, their own internal systems were still using older technologies and worked on processes and systems that had largely remained untouched for the past decade.

Through the DevOps and Agile movements, that transformational change had already transformed the way applications were being built, tested, and deployed, and most of the tasks in the development value stream were automated and integrated, resulting in organizations moving up to continuous delivery and continuous deployment. Similar to the transformation seen with DevOps, with AIOps gaining traction, we are going to witness a transformational change that will drastically change the way IT operations has been run in the past. Old processes, systems, tools, and ways of working will give way to the AIOps way of operating and help realize the vision of NoOps, where operations work seamlessly without disruption in an automated way without human supervision and intervention.

Enterprises today are at various levels of maturity when it comes to automation. Most are yet to achieve a high level of maturity in automation, partly because changing processes, breaking down walls of organization structures, and deploying new technology are complex, time-consuming tasks. In some organizations, where digitization and cloud computing programs are a part of large-scale transformation; AIOps and associated technology and process changes are enabling a complete transformation of digitization. One key factor in adopting AIOps has been a lack of comprehensive process and technology guidance in this domain. There is limited guidance available, and most of it is focused on products that vendors are trying to sell as a one-stop solution, which will leapfrog an

PREFACE

organization to the next level. With this publication we are aiming to provide hands-on pragmatic guidance on how an organization can adopt these changes in the origination and what the pitfalls are to be avoided. Implementing AIOps is not about deploying an event correlation tool; it is about infusing AI and algorithms and automating all aspects of operations.

Most organizations today are at a level where they have automated individual tasks and have automated certain processes. Some have also achieved automation of complex end-to-end processes. However, an end-to-end self-sustaining model for automation is still missing. Automating a few processes versus automating everything that can be automated is a change that people haven't yet initiated. Automating a few processes and then scaling the automation across teams, functions, and departments is still work in progress in most organizations.

AIOps is not just a technology or process change; it is a cultural change where humans and AI work together in tandem supporting and augmenting each other's capabilities to achieve efficiencies and scalability, which IT teams have only wished for but have been unable to achieve. The promise of NoOps has failed to materialize because organizations have not been able to take a holistic and end-to-end approach to AIOps and instead have relied on quick fixes in the form of siloed tools that are deployed.

In enterprises that have moved up the maturity curve in automation, most activities that can be automated fully or partially are automated, and AI drives decision-making rather than humans using a dashboard and analytics to arrive at the next steps. It is a mindset change where humans have to accept the supremacy of the AI systems in certain areas and let go of the control mindset inherent in our species. IT teams are used to certainty and rule-based systems that have dominated the IT landscape for decades. AI-based systems are probabilistic and not definitive; thus, some enterprises have deployed AIOps systems but configured them in a manner where they are nothing but rule-based systems. Accepting probabilistic systems rather than step-by-step rule-based systems is a mindset change, and one needs to accept the risks that come with

implementing such systems and controls to ensure best practices for running AI-based systems; human control on some of the actions is incorporated as part of the process changes.

Enterprises are also evolving to change their operating models to align with Agile principles; they are starting to integrate applications and infrastructure teams and are adopting new practices such as site reliability engineering to make operations support highly reliable and high-availability mission-critical systems and Internet-scale applications. AIOps is a key pillar in site reliability engineering-based operations. The book *Hands On Guide to AgileOps* published in 2021 by Apress provides hands-on guidance on how to adopt Agile for IT operations.

This book will take you through all aspects of AIOps including AI and machine learning algorithms and demonstrate how some of the features in AIOps are implemented in enterprise AIOps platforms. You will be able to leverage the capabilities provided by AIOps processes and platforms and take your organization to a higher level of maturity in operations. We will continue to provide updated content and guidance on the companion website, www.AgileInfraOps.com, where you will find articles and best practices as processes and technology continues to evolve.

Introduction

This book is a hands-on guide to understanding AIOps in detail and assisting in implementing its technologies and processes in the organization.

The book explains the IT industry's need for AIOps, its architecture, and the technologies that contributed to its evolution. Readers will be able to grasp the core theoretical concepts around AIOps as well as go deeper by implementing hands-on examples and exercises that leverage machine learning techniques to implement AIOps. The book also provides guidance on how to set up AIOps in an enterprise and what pitfalls to avoid to complete a successful implementation. The book also explains the role of AIOps in the SRE and DevOps model and provides a detailed explanation of the enablement of key SRE principles by AIOps.

This hands-on book also provides an implementation of multiple AIOps use cases and provides sample code that you can run to better grasp some of the underlying principles and technologies that form the core of AIOps.

This book shares the best practices, processes, and guidelines to establish AIOps practice and systems in enterprises and the methods to measure the outcomes and continuously evolve.

CHAPTER 1

What Is AIOps?

This chapter introduces artificial intelligence for IT operations (abbreviated as AIOps). In today's rapidly transforming application and infrastructure landscape and adoption of cloud-native technologies, organizations are finding it difficult to provide 24/7 operations that can scale and meet the needs of businesses that now want much higher availability and agility to change based on customer and market feedback. This chapter also provides details on the benefits that AIOps brings to the table and how it supports the digitization journey of enterprises.

Introduction to AIOps

AIOps is a buzzword in the operations world and was coined by Gartner in 2016. As mentioned, it means implementing artificial intelligence for IT operations. AIOps refers to a transformational approach to running operations using AI and machine learning technologies in various operations domains such as monitoring, observability, event correlation, service management, and automation. With the exponential growth seen in application and platform diversity, including the movement to microservices and cloud architectures, there is an enormous amount of data being generated in operations. The operations teams are overwhelmed with this vast amount of data and the diversity in applications, platforms, and infrastructures in the environment. Most enterprises today are rapidly migrating and adopting new technologies such as cloud and microservices architecture, and thus the rate of

change in infrastructure and platforms is unlike anything seen before. The challenge in IT operations is to run steady-state operations without disruption and also support this agility and migration and bring new services into operations. These disruptions and changes are putting an enormous strain on the operations teams. Processes and systems that have worked in the past are not working anymore, and the new digitalized world with rapid changes both in applications and infrastructure is resulting in newer challenges. Thus, AIOps has evolved over the last few years as a potential solution to the operational challenges of the new model.

The huge amount of data getting generated from monitoring and observability systems is one of the sources of data that is fed into AIOps-based systems, and then AI and machine learning techniques are used to make sense of the data and filter the noise from critical events. This results in automating most of the tasks that were manual before and that relied on human judgment and tribal knowledge. The events that can cause disruption in business operations and are the root cause are efficiently identified using analytics techniques and thus provide immediate notification to the groups that are resolving them. Without AIOps this process will be difficult to run with changes in technology happening at a rapid pace; relying on older systems and tribal knowledge would mean operations is not scalable and predictable.

All of this is enabled by the emergence and maturity of artificial intelligence and machine learning technologies, which are the foundation of AIOps.

Artificial intelligence has transformed the way systems are developed and business processes are run. AI is everywhere from the image processing in your phone to recommendation engines on Amazon that provide you with new product recommendations based on your preferences. Face recognition and image beautification on phones are examples of applications that people are consuming every day without even knowing that artificial intelligence is powering these applications. Natural language processing advancements have transformed the way

we interact with applications. Today voice assistants like Alexa, Siri, and Cortana are changing the way we communicate with content.

Information technology has leveraged these technologies to solve varying business problems in the areas of building recommendation systems, predictive systems, image recognition, voice recognition, text extraction, and natural language understanding systems.

However, when it comes to solving IT problems using artificial intelligence, enterprises and technology companies have yet to embrace this fully.

Finally, the AI technologies that IT teams have used to deliver exciting new applications for consumers and businesses are now finding their way into monitoring and managing the IT technologies. Thus, a new class of systems that are using algorithms to run IT operations was born.

AIOps is a term Gartner invented to describe a general trend of applying AI techniques to IT operations data sources to provide additional insights. AIOps is essentially a feature or set of features to analyze, combine, and collect data.

According to Garner, “By 2023, 40% of DevOps teams will augment application and infrastructure monitoring tools with artificial intelligence for IT operations (AIOps) platform capabilities.” AIOps platforms are platforms that “utilize big data, modern machine learning and other advanced analytics technologies to directly and indirectly enhance IT operations functions with proactive, personal and dynamic insight.”

Figure 1-1 defines the various areas in IT operations that are included in AIOps. These include monitoring, event analytics, predictive and recommendation systems, collaboration and engagement, and reporting and dashboarding technologies.

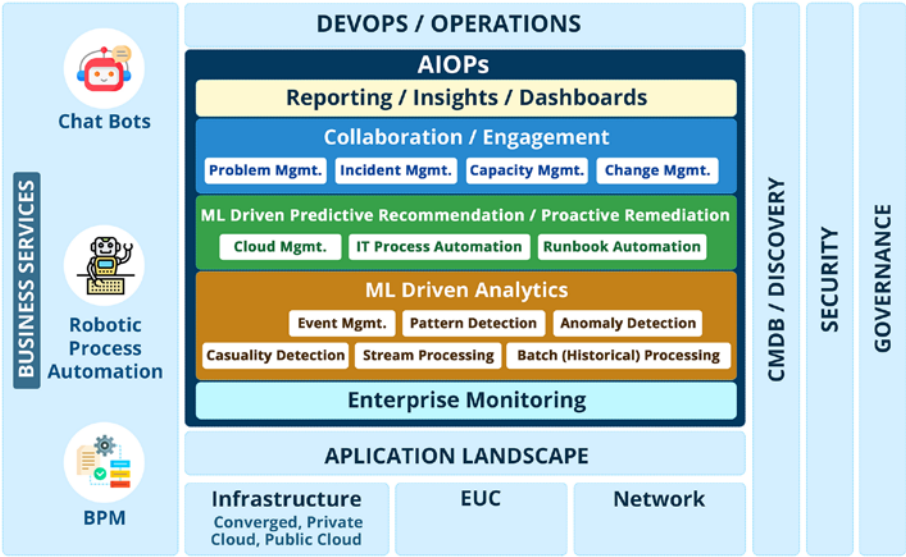


Figure 1-1. AIOps in the IT operations landscape

AIOps covers various layers of information technology. From the network to the endpoints, everything in IT can use AIOps technologies to reap the benefits that AIOps provides.

Enterprise monitoring provides a real-time data feed to the AIOps system to perform ML-driven correlation and analysis using various techniques to detect patterns and anomalies, as well as perform causal-impact analysis. This is one of the most important stages as this analysis needs to consider both real-time streaming data as well as historical data to provide predictive recommendations or proactive remediations, which then get executed using IT process or runbook automation tools. The recommendations for resolution help the organizations in achieving end-to-end automation by resolving problems without human intervention.

The reporting and dashboard layer provides views for different IT teams and stakeholders to collaborate and manage incidents, capacity, change, and problems to further support business by providing KPIs and SLAs that are driven by insights and provide an element of predictive analytics to make operations more proactive.

AIOps systems leverage a Configuration Management Database (CMDB) to improve the quality of correlations and accuracy of predictions and recommendations, but organizations usually struggle in maintaining the accuracy of the CMDB and discovery data due to the ever-changing infrastructure landscape. With cloud computing, it is practically impossible to update the CMDB using traditional tools and processes. An AIOps system solves this problem by automatically populating the required missing data in the CMDB. This well-oiled engine of AIOps has to work within organizations' security policies defined under their governance framework. Various compliance needs like GDPR, data classifications, etc., should be considered at each layer of the AIOps engine while they are being integrated or set up. Once AIOps systems get up and running, they learn patterns, anomalies, and behavior using data over a period of time. Gradually, based on maturity, the AIOps system gets consumed by other technology or business units such as ChatOps, robotic process automation, business process automation, etc. More complex business process workflow or chat responses can be triggered based on the AIOps system recommendations.

Just like in software engineering, continuous integration and continuous deployment integrates different activities of development, testing, and deployment of applications and shares feedback for improvement. Similarly, AIOps is something that provides seamless integrations between various operations components and provides feedback for continuous service improvement.