LogicMonitor

AIOps for Monitoring

Key concepts, benefits, and AIOps use cases for your current and future monitoring strategy

Table of contents

Part 1: What is AIOps for monitoring?

3 Introduction

4 Chapter 1: Introduction to AIOps
- 5 Keeping pace with evolving IT environments
- 5 How IT is changing
- 6 Benefits for IT teams and leaders
- 6 IT Execs are strategic and outcome-driven
- 7 The role of observability

8 Chapter 2: AIOps fundamentals
- 8 Using AIOps to drive better business outcomes
- 10 Start with an early warning system
- 12 Getting started

13 Chapter 3: Implementing your AIOps initiative
- 14 Solve for complexity with AIOps
- 14 AIOps processes and platforms
- 14 Solving the big data problem
- 14 Ensuring meaningful alerts
- 14 The importance of a single platform solution
- 15 Getting started with AIOps

Part 2: How to use AIOps for monitoring

16 Introduction

17 Chapter 4: Early warning system overview & benefits
- 18 Surfacing hidden signals to power rapid action
- 18 Preventing downtime
- 18 Minimizing alert fatigue
- 19 Opening up visibility, context, and understanding
- 19 Preventing widespread business problems instead of reacting to them

20 Chapter 5: Dynamic thresholds
- 20 Discover the exceptions with dynamic thresholds
- 20 Identifying patterns and pattern types
- 22 Benefits that span the IT team
- 22 Benefits for IT teams
- 22 Benefits for IT leaders
- 22 Conclusion

23 Chapter 6: Root cause analysis
- 23 Identifying the true source of an issue
- 24 Using topology mapping to find the root cause
- 25 How AIOps-based root cause analysis works
- 26 Configuring root cause analysis
- 27 Benefits for IT teams

28 Chapter 7: Forecasting
- 28 Demand and capacity planning data forecasting
- 29 Viewing and understanding forecast graphs
- 29 Streamlining the planning process

Part 3: The future of AIOps

30 Introduction

31 Chapter 8: Putting AIOps use cases into action
- 31 Empowering IT operations with the big picture
- 32 Enabling managed service providers to enhance the customer experience
- 32 Breaking down silos and sharing global insights
- 33 Driving innovation and agility with DevOps
- 33 Automatically choosing the right anomaly detection algorithm
- 33 Gaining DevOps observability with machine learning anomaly detection
- 34 Extending the outlook beyond fixing issues

34 Chapter 9: Future for AIOps: Automated issue remediation
- 35 Driving optimization with a failure prevention system
- 35 Defining and automating action
- 35 Nonstop learning about environments
- 36 Tracking the next wave of AIOps
- 36 Opening up the potential of business data
- 37 Glossary

Part 1

What is AIOps for monitoring?

Key concepts for IT organizations considering artificial intelligence

Introduction

Technology innovations transform human behavior permanently. In the past decade, we have improved as a society by embracing digital lives that drive faster collaboration and automation and save us a significant amount of time. The IT Operations landscape is not any different, and artificial intelligence (AI) is at the forefront of that. The goal is to make a step function increase in productivity by adding more automation so that IT organizations can shift focus to solving important and difficult problems that cannot be easily automated. Every year, more of that mundane work will continue to get automated by AI.

In this book, we will look past the hype, and take a closer look at what AIOps actually is —and what it can help you achieve. First, we'll explore why AIOps is necessary. We'll talk about how environments, and the teams that manage them, have evolved in an increasingly multi-cloud world. We'll explore how observability of the infrastructure is fundamental for helping IT organizations work more proactively and strategically.

Next, we'll present a basic definition of AIOps, and its ability to apply an automated early warning system to discover potential issues and provide the context you need to take action. By knowing what is happening, where it is happening, and why it is happening, you can solve problems faster, deliver the best possible service quality, and stay focused on driving digital transformation. This is a fundamental first step toward a failure prevention system, or automated remediation.

Finally, we will talk about the primary steps involved in creating your own AIOps initiative. We'll take a closer look at how AIOps helps organizations bring together big data from across the technology stack, surface only the most relevant issues, and apply advanced algorithms to deliver more meaningful alerts. Our book will conclude with an overview of some key best practices you can apply to implement the most effective AIOps solution, and realize the best return on your investment.

Chapter 1: Introduction to AIOps

Now more than ever, IT organizations are playing a more strategic role in driving business outcomes (revenue growth, cost optimization, risk mitigation, and operational efficiency) and accelerating innovation. All eyes are on IT departments, which are under increasing pressure to work faster and deliver results while reducing costs. This becomes challenging as infrastructures become more complex, diverse, and dynamic.

Chances are, you're grappling with a mix of on-premises, private cloud, different hosting providers. When you consider additional environments like public cloud, IaaS, PaaS, new networking technologies, and infrastructure as code provision, you've got multiple suppliers to manage, and multiple moving parts.

In this chapter

In this chapter, we'll take a brief look at the new challenges that today's infrastructures, applications, and workloads present, and explore what they mean to IT teams, leaders, and executives. We'll talk about the changing expectations for IT organizations, and how they are embracing a more strategic role in driving digital transformation. And we'll show how an AIOps approach that builds on observability and automation enables IT to gain the insights they need to work more proactively, and perform more effectively.

In this chapter, you will learn:

- How expectations are changing for IT organizations
- The role IT plays in driving digital transformation
- How AIOps and observability enable a proactive it organization

Keeping pace with evolving IT environments

Not so long ago, the IT infrastructure was an environment that could be seen, understood, and managed. However, today's enterprises are embracing digital transformation and modernizing their infrastructures to deliver products faster, meet new customer expectations, and stay ahead of the competition. It's not unusual for modern hybrid infrastructures to have a mix of resources in the cloud, and others running on-premises in physical data centers. In a multi-cloud world, network environments are becoming more complicated, complex, and opaque.

According to the Nutanix Enterprise Cloud Index Report, the majority of enterprises plan to shift to hybrid cloud architectures with 86% of respondents selecting hybrid cloud as their "ideal IT operating model."[1]

It's clear that new customer wants and needs are driving digital transformation, including a rapid evolution of modern applications and workloads. Traditional data center strategies can no longer support them, and it would be cost prohibitive and too time consuming for businesses to try to modernize their data centers in this way.

According to Gartner, cloud application services (SaaS) will grow to $195.2 billion in 2023, and the second-largest market segment in cloud system infrastructure services, or infrastructure as a service (IaaS), will reach $150.3 billion in 2023.[2]

How IT is changing

As IT infrastructure environments are changing, the role of IT is evolving as well. After years of being considered a service organization, IT has to become more strategic and creative, and must play a leading role. IT Ops professionals are no longer focusing on their old mission of "keeping the lights on" by keeping existing internal systems up and running. Instead, they are increasingly focusing on providing the resources that line of business needs to power high-profile business outcomes.

In a recent survey, 41% of technology officers reported digital transformation as their top strategic priority.[3]

Today's leading technologists drive innovation and external services such as E-commerce, mobile offerings, and new IoT-based services. Salesforce reports that customer expectations are at an all-time high. Their survey found that 67% of consumers believe that their standards for good experiences are higher than ever.[4]

The bottom line is that today's IT infrastructure plays a key part in driving digital business success. The stakes are high, and both the employee and customer experience (CX) are entirely dependent on the underlying infrastructure working and performing well.

[1] Enterprise Cloud Index, 4th Annual [Nutanix], 2022

[2] Gartner Forecasts Worldwide Public Cloud End-User Spending to Reach Nearly $600 Billion in 2023 [Gartner], Oct. 2022

[3] Technology Officers Pulse Survey 2019 [Korn Ferry], 2019

[4] What are Customer Expectations, and How Have They Changed? [Salesforce], n.d.

Benefits for IT teams and leaders

Together, new responsibilities and evolving environments are bringing new pressures to bear on IT professionals at every level. Whether they are working on the front lines as IT Operations professionals, developers, or members of support teams, or setting strategic priorities as executives, IT professionals need improved observability into their infrastructures, so they can make better, faster decisions.

Benefits for IT teams

For people on IT teams such as IT analysts, system engineers, and infrastructure architects, effective execution is top of mind. If you're on an IT team, it's likely you're routinely building dashboards, managing alerts, and responding to issues using monitoring solutions, but are also striving to work more strategically. Ideally, you would probably prefer to spend less time fighting fires, and more time focusing on larger-scale projects.

Today's increasingly distributed and complex IT infrastructure and application architectures can stand in your way. Scattered across on-|prem, private, and public cloud environments, they make ensuring performance and availability increasingly difficult.

IT teams are also under pressure to get things done faster and better, as the cadence of releases and changes accelerates. Detecting a network issue after it has already impacted a business process or customer experience is simply too late. That means you can find yourself in an endless cycle of fixing urgent issues.

Workload concerns are the highest they've been in the history of our IT Skills and Salary Report, and most respondents agreed that automation may be one solution to reduce time-consuming tasks that are not high priority.[5]

By freeing you up and giving you more time to operate proactively, you could gain opportunities to tackle more interesting projects, and focus on driving business outcomes. Playing a more strategic role will not only keep you more satisfied, but help move up to higher roles and raise your professional profile.

Benefits for IT leaders

If you are an IT leader like a manager or director, aligning technology with your organization's desired business outcomes is critical. Every day you're working closely with senior leadership to develop and refine initiatives that support the organization's top business priorities. Caught between a rock and a hard place, you are expected to ensure maximum uptime and performance for your IT infrastructure, even as you are asked to transform CX through major initiatives.

As you manage infrastructures spread over diverse cloud and on-premises environments, finding and hiring the right IT support professionals with the right skills is not easy. You need to get out in front of these challenges by gaining as much insight into your environments as you can, so you can make better strategic decisions.

The ability to gain visibility into what is happening on enterprise networks has once again topped the agenda for network managers and buyers, according to the latest Computer Weekly/TechTarget IT Priorities survey.[6]

If you are an IT leader, you are also looking for better ways to de-risk operations and gain control as you evolve IT. With the right monitoring solution for visibility and insight, you can accelerate mean time to resolution (MTTR) as you adopt next-generation technologies, and deliver on your organization's corporate strategy.

IT execs are strategic and outcome-driven

If you are a member of the executive team, you are looking at the big picture. For a CIO, CTO, or other senior executive, technology is the engine that supports their company's long-term corporate goals and priorities. You understand that the relationship between IT and business processes is becoming more connected all the time; however, although you are increasingly adopting cloud-based solutions to drive innovation, these diverse environments introduce new levels of complexity that can slow progress and drive costs up.

> ***"By 2024, 25% of traditional large enterprise CIOs will be held accountable for digital business operational results, effectively becoming 'COO by proxy.'"[7]***

IT executives like you are looking for deeper visibility into infrastructure health to understand how it is impacting applications, IT services and business processes directly, in order to help the entire organization become more data-driven. With the right context and insights, you can help your colleagues in the c-suite see around the corner—and prepare for what's next.

[5] 12 Challenges Facing IT Professionals [Global Knowledge], Mar. 2020

[6] Scroxton, Alex. IT Priorities 2019: Security Concerns Top Agenda For Network Buyers Again [ComputerWeekly.com], Feb. 2019

[7] Top Priorities for IT: Leadership Vision for 2021 [Gartner], 2020

The role of observability

It's clear that today's IT organizations have more to observe and more to understand in their environments. To see more, know more, and do more, you need full insight into your systems, workloads, and processes, including metrics, logs, and tracing. That level of insight must extend across your dispersed environments and infrastructure, whether they reside on-premises, in the cloud, or within microservices. Most importantly, you need to achieve this level of observability in a timely way, using it to detect performance issues before they escalate into business issues.

Observability is a subset of AIOps (Artificial Intelligence for IT operations) and monitoring. Simply put, observability is about ensuring that system data is 'observable' from an availability and performance perspective by a monitoring platform. In recent years there has been the emergence of Observability platforms, or the ability to monitor metric, application and log data in a single platform approach. Recently Unified Observability has begun to be important, as it's not only about monitoring these data sets, but ensuring that they are correlated and in context of each other. AIOps capabilities are then applied to automatically filter out the noise from the streams of data, to enable early warning of issues before widespread business impact and to proactively avoid failures from happening.

AIOps builds on observability and automation to enable you to support increasingly complex environments. An effective AIOps approach should work seamlessly with existing data sources, and enable IT to be more proactive and strategic.

In the chapters ahead, we'll provide a full definition of AIOps, and an overview of some of the key features and components that it utilizes. We'll show you how it can impact your business, and enable new levels of service quality and availability, to power better business outcomes. And we'll provide an overview of some of the steps you need to develop, deploy, and sustain your own AIOps initiative.

Observability is a subset of AIOps and monitoring. Simply put, observability is about ensuring that system data is 'observable' from an availability and performance perspective by a monitoring platform.

In recent years there has been the emergence of Observability platforms, or the ability to monitor metric, application and log data in a single platform approach. Recently Unified Observability has begun to be important, as it's not only about monitoring these data sets, but ensuring that they are correlated and in context of each other. AIOps capabilities are then applied to automatically filter out the noise from the data, to enable early warning of issues before widespread business impact, which proactively prevents failures such as outages.

Chapter 2: AIOps fundamentals

It's clear that IT teams are grappling with more data than ever before, along with a growing mix of products that they depend on to help them monitor that data. Unless they take proactive steps to capture, understand, and act on this flood of new data, the consequences could be severe, ranging from performance issues and outages to a wider failure to drive digital transformation effectively.

In this chapter

In our previous chapter, we talked about some of the trends and challenges driving the need for AIOps, and discussed why comprehensive observability is fundamental to enabling it. In Chapter 2, we will present a general overview of some of the fundamentals of AIOps and its key building blocks, including anomaly detection, dynamic thresholds, root cause analysis, and forecasting. And we'll discuss why AIOps has become such a compelling solution for today's increasingly diverse, fast-moving environments.

In this chapter, you will learn:

- The AIOps fundamentals
- Key building blocks required for an AIOps initiative
- Why AIOps is the best solution for IT complexity

Using AIOps to drive better business outcomes

AIOps is all about the application of ML/AI algorithms that are able to automatically detect anomalies, such as change or capacity issues in an infrastructure, application, or service, before they become problems. It starts with observability as a crucial first step—being able to observe metric, log and application data in context via a monitoring platform. Automation plays a key role in enabling AIOps to find anomalies, build insights, and respond to them. An early warning system is fundamental to driving AIOps, by showing what is happening, where it is happening, and why it is happening.

AIOps is not only about issue mitigation, but also about continuous optimization. Its capabilities can provide support for a failure prevention system that continually learns about the environment, to optimize the technology stack so that it drives business outcomes. With an AIOps approach, IT professionals can free themselves from focusing on baseline "keeping the lights on" responsibilities and assume a more strategic in driving innovation and digital transformation.

AIOps was created specifically to address the challenges of today's modern hybrid infrastructures, and enable businesses to see what's coming, spend less time troubleshooting, and more time innovating. It lets organizations break the endless cycle of reactive monitoring to embrace observability and get out in front of infrastructure issues, moving from a proactive to a predictive approach before they can impact business operations and the customer experience.

Its key capabilities include:

- Ingesting data from multiple sources, agnostic to source or vendor, to provide a comprehensive picture of IT environments
- Performing real-time analysis at the point of ingestion to open up faster insights and enable more proactive responses to issues
- Performing historical analysis of stored data to better inform insights with additional context
- Leveraging machine learning to surface the most relevant alerts, and suppress redundant data that can distract and slow IT responsiveness
- Employing automation to initiate an action or next step based on insights and analytics, to enable rapid responses to issues before they impact the business

AIOps is not only required to help IT deal with increasing technical complexity, but also to ensure the need to support business outcomes efficiently and effectively through technology is satisfied. It helps IT Ops teams handle their primary responsibilities more effectively while enabling them to better align their activities with key business stakeholders.

> *“By embracing AI augmented automation, IT teams can better learn the skills of AI and position themselves to have more effective partnerships with peripheral business units. In fact, by 2023, 40% of infrastructure and operations (I&O) teams will use AI-augmented automation in large enterprises, resulting in higher IT productivity with greater agility and scalability.”*[8]

AIOps takes two approaches. It can be domain agnostic, integrating with a variety of different services to collect data, or domain-centric, primarily collecting the data that is required on its own. In this chapter, our focus will be on domain-centric AIOps.

Maximizing availability and improving collaboration

What happens when you empower IT Ops to do their jobs more proactively—and bring them closer to business imperatives? By providing the relevant data that ops teams need to make better decisions, faster, you unleash service improvements that drive compelling business outcomes.

Maximizing service quality and availability

AIOps gives IT teams not only the visibility and insight they require to spot potential issues, but the context and automation they need to prevent incidents. For example, instead of waiting for an outage to occur, AIOps lets you identify troublesome resource usage patterns that could cause an outage if they continue. After discovering the potential problem, AIOps could deliver actionable recommendations to avoid the outage. When you integrate AIOps with an automation platform, you can address the issue even more quickly, executing recommendations and fixing resource usage well before a problem can escalate and impact your users.

AIOps can also give you the insight you need to support more stable, highly available customer-facing services. You can determine which applications and infrastructure issues are most likely to impact the stability of your environment—and the user experience—and take steps to prioritize and eliminate them. The result is a superior experience, improved user satisfaction, better long-term customer retention, and enhanced revenues.

Improving collaboration to assess problems

As data becomes more essential to every organization, business and technical teams are becoming more collaborative. Technical teams are looking to align better with business imperatives, while business stakeholders seek metrics to understand how technical issues affect their priorities.

AIOps capabilities let you apply contextualized data to create more precise, tangible insight into how an incident will impact your operations, and your end customers. Instead of collecting and analyzing metric, log, and application data from several systems, infrastructure, operations, DevOps teams, and other business stakeholders can easily access and visualize insights and get answers in the context of their role.

AIOps is ready for prime time

With its tremendous potential to power better business outcomes, it’s not surprising that enterprise organizations across a variety of industries are embracing AIOps. According to Gartner, “**by 2023, 40% of DevOps teams will augment application and infrastructure monitoring tools with AIOps platform capabilities.**”[9]

One of the reasons AIOps is enjoying such enthusiastic adoption is that its capabilities track closely with the changing roles of today’s technology teams. More than ever, I&O teams are becoming more strategic. They are focusing less on managing technology like data centers, colocation and the cloud, and more on how they can drive growth, spark innovation, and support and enable their organization’s business strategy.

> *“AIOps platforms address I&O leaders’ need for operations support by combining big data and machine learning functionality to analyze the ever-increasing volume, variety and velocity of data generated by IT in response to digital transformation.”*[10]

[8] Gartner Predicts the Future of AI Technologies [Gartner], Feb. 2020

[9] Market Guide for AIOps Platforms [Gartner], Nov. 2019

[10] Market Guide for AIOps Platforms [Gartner], Nov. 2019

For example, an organization might be experiencing a sudden slowdown in database performance. AIOps could provide insights that reveal a correlated increase in connections from another service. It even could go a step further, and note that the troublesome service had been reconfigured just before the database issues had begun. It quickly becomes clear that a configuration error is the root cause of the database issue. To resolve it, IT can simply revert the service to its previous configuration, and restore the database to its normal performance.

Start with an early warning system

When the infrastructure is the engine that drives so many of your key business operations, the stakes are incredibly high. It's not enough to simply see and fix problems quickly. To succeed, you need the ability to see what's coming and respond to it before it happens—and spot potential issues even when they are far on the horizon.

An effective AIOps approach will act as an early warning system that alerts you to the most relevant issues in your environment. It will not only spot issues, but detect the warning signs and symptoms that precede them, like patterns or anomalies in alerts or performance data—and warn users about them. These early warnings can trigger actions, such as integrations and custom scripts, to prevent issue occurrence.

By alerting users sooner, an early warning system helps you prevent outages, save time and money, and maximize the availability of your most important business and customer services. It also helps you work smarter, enabling your IT ops team to zero in on resources that cause outages and speed up your mean time to repair (MTTR).

Building blocks of the early warning system

To deliver these smart, proactive capabilities, an AIOps early warning system is key. Essentially, an early warning system helps you understand what is happening, where it is happening and why it is happening. It requires four key building blocks, including anomaly detection, dynamic thresholds, root cause analysis (RCA), and forecasting capabilities. Anomaly detection uses powerful ML algorithms to spot unusual activities that could impact your infrastructure health, and support the dynamic thresholds that alert you about issues. If a problem emerges, RCA applies automation to help you trace its source. And data forecasting helps you identify trends that can keep issues from emerging in the future. Let's take a closer look at each capability, and how they work together in an early warning system.

Enterprise AIOps platform

An enterprise AIOps platform is composed of three key features: full observability, an early warning system, and a failure prevention system.

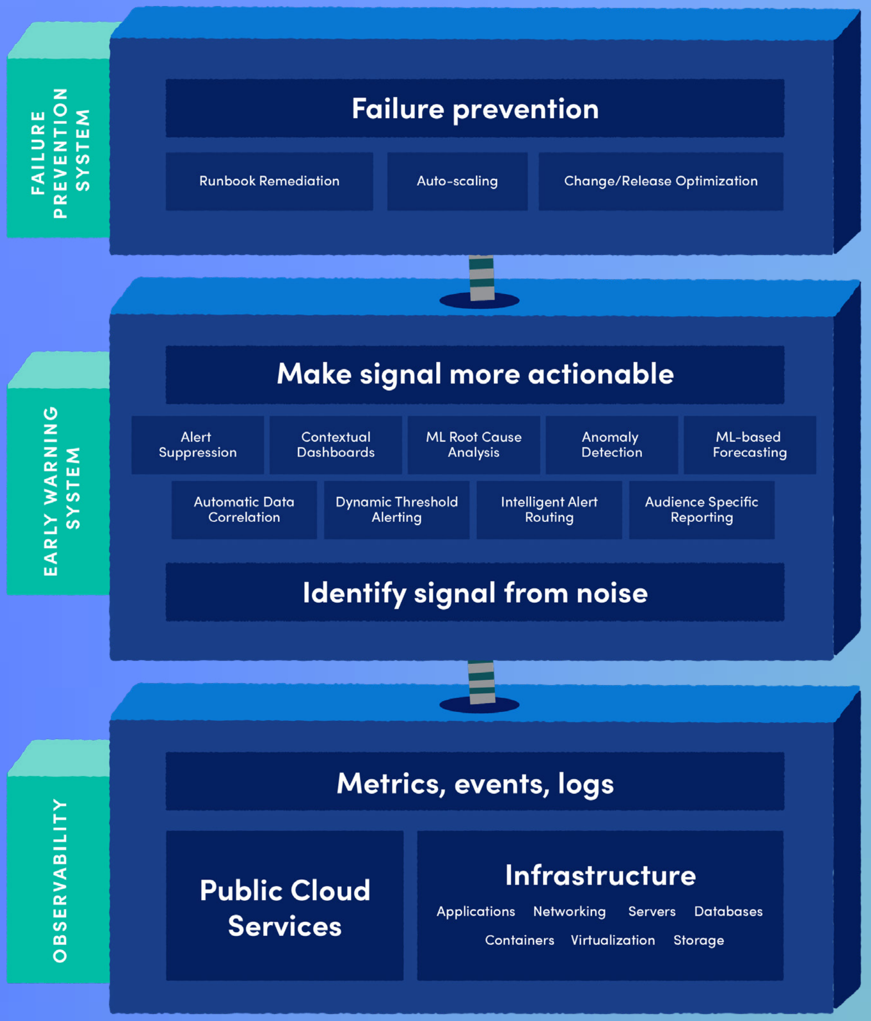

Anomaly detection

Anomaly detection is what powers the dynamic threshold capabilities in an AIOps early warning system. It gives you a holistic view into the health of the resources you're monitoring, along with the insight you need to dig deeper and mitigate critical events that could impact your business. With anomaly detection, you can see deviations that occur within your resources, and compare these anomalies to key historical signals. Anomaly detection also utilizes log information to identify anomalies, correlate them with metrics from monitored resources, and surface them to provide additional context and insights for alerts.

Anomaly detection gives you more than just visibility. It provides an extra layer of intelligence that can help you understand the expected performance of your resources and know when actual performance is different from those expectations. For example, you could graph an overview comparing a resource's actual bandwidth utilization to its maximum average trend and highlight any unusual events. Anomaly detection makes it easy to not only better understand your resource health, but also troubleshoot faster and more effectively.

Dynamic thresholds

Dynamic thresholds depend on anomaly detection, because as infrastructures grow in complexity, they are generating more metrics and data than ever. Keeping track of the most relevant issues and events is impossible to do manually, especially in large environments. With AIOps, you can set up dynamic alerts that are only sent for anomalies. This reduces alert noise and fatigue, while helping to eliminate the need for manual threshold tuning.

Dynamic thresholds use anomaly detection algorithms to detect a resource's expected range, based on its past performance. Instead of sharing everything with IT Ops, they issue alert notifications only when they spot something that lies outside of this normal range. Dynamic thresholds will catch anomalies in metric values, metric rate of change, and even metric patterns, like a drop in traffic where it isn't normal. In addition to ensuring that alerts are generated for anomalies, dynamic thresholds can be used to reduce noise where static thresholds aren't tuned well.

Dynamic thresholds require a minimum set of hours of training data. As more data is collected, the algorithm is continuously refined, and can use several days' worth of recent historical data to inform its expected data range calculations. And as time passes, dynamic thresholds can continue to learn from historical data, and become more accurate. As accuracy increases, you can ensure your team is focusing on what's really important, so they can work more efficiently and spot potential issues sooner.

For example, an organization could use dynamic thresholds to support seasonal, scheduled events and deliver a superior customer experience. A retailer might run a large sale on Black Friday, and utilize microtransaction monitoring logs to ensure that its ecommerce platforms are responsive, and keeping pace with a corresponding flood of new transactions.

Building on knowledge from previous data, dynamic thresholds could continue to spot anomalies and alert IT Operations teams about issues, even when thresholds have dramatically risen.

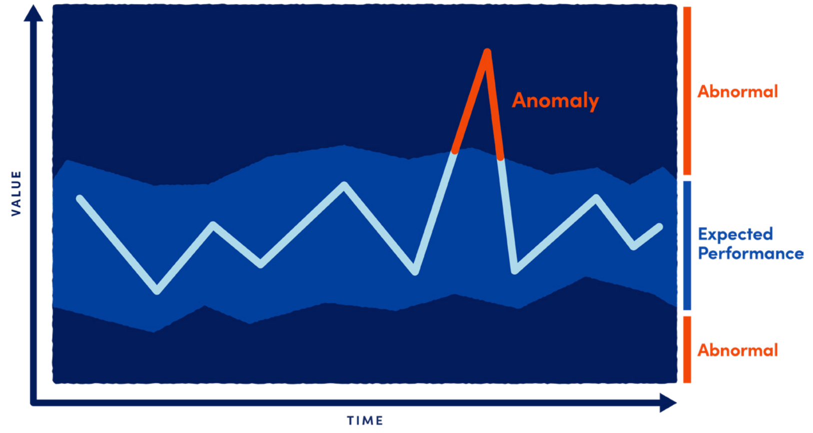

This graph identifies the anomalous (abnormal) action in red because it exceeds the expected performance range.

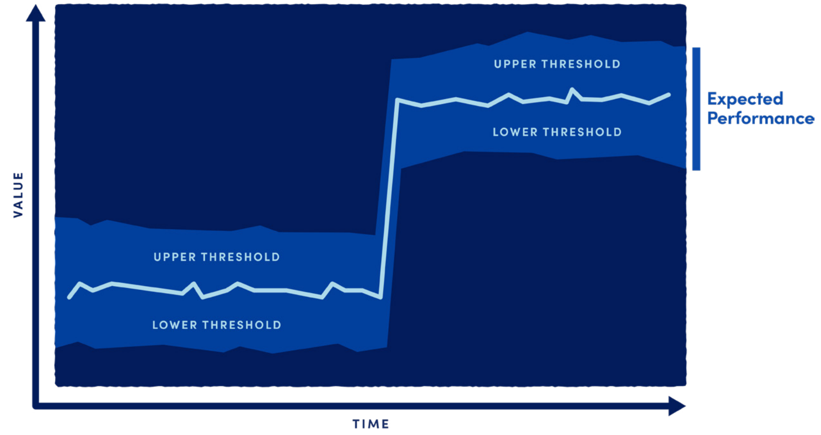

The expected range, or dynamic threshold, is represented in light blue. This range has an upper and lower limit representing normal or expected behavior, anything outside of this range will be considered abnormal or anomalous.

Root cause analysis

It's one thing to discover an issue early on, but to be truly proactive in resolving it, you need to be able to look more closely and quickly pinpoint the origin of a problem. Root Cause Analysis uses automation to dive deeper, and requires a deep understanding of relationships and dependencies within your network, informed by topology mapping. Using topology mapping, it examines automatically discovered relationships between your monitored resources to identify the root cause for triggered alerts. Then it notifies your users of the originating issue. RCA can also give you the option to suppress notification routing for alerts that are dependent on the originating alert. This lets you dramatically reduce alert noise for events in which a parent resource has gone down or becomes unreachable and has caused its dependent resources to go into alert as well.

RCA is also about putting different data sources in context with one another, such as correlating metric alert data with the right log events. For IT Ops, RCA accelerates mean time to resolution by issuing alert notifications that clearly identify an issue's root cause.

With the right AIOps platform, you can also utilize filterable in-app alert data, to enable your team to precisely zero in on resources that play a key role in outages, so they can identify and resolve issues faster. RCA also helps you minimize alert fatigue. Since your IT team is only notified about the root cause issue, they can focus on fixing what's important, instead of getting overwhelmed by dependent side effects.

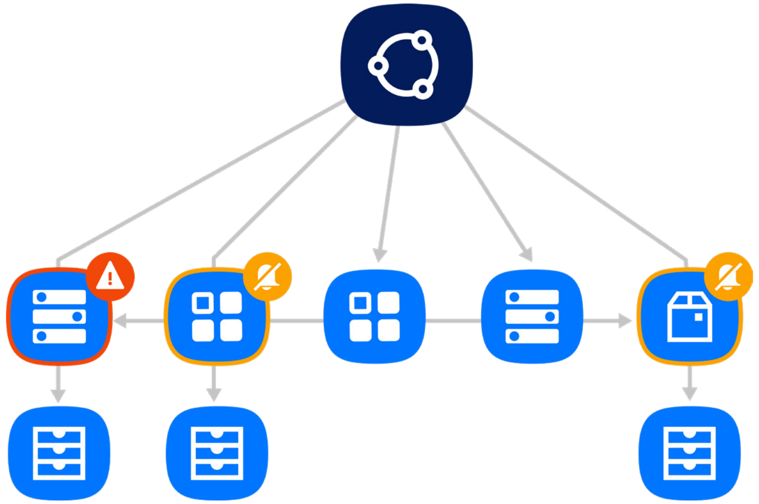

With root cause analysis users can see their environments visually represented with a topology map. The map reflects what devices are in alert (red error) and the dependent devices where alerts have been silenced (yellow bell).

Forecasting

Accelerating issue resolution is a compelling benefit of AIOps, but the power of AI really shines when you can use it to get ahead of the curve. Data forecasting lets you predict future trends for your monitored infrastructure, using past performance as the basis. With the right AIOps capabilities, you can identify and remove anomalies and missing data from the set of data you choose for forecasting. Then the platform applies a capacity trending algorithm to the sample to find a model that best fits the data you have collected. The platform can then calculate future data based on these model parameters. When you combine data forecasting with your alerting system, it becomes even more powerful. You can determine whether an alert represents a one-time anomaly, requires immediate attention, or will require attention in the near future.

Forecasting can help you identify upcoming issues before they trigger alerts, so you can prevent downtime. It's also a great tool for budget planning and resource management, by predicting the health and performance of your monitored devices. When you track infrastructure components with specific lifetimes or capacities, forecasting can give you insight into the timeframe and magnitude of recurring events, as well as upcoming expenses.

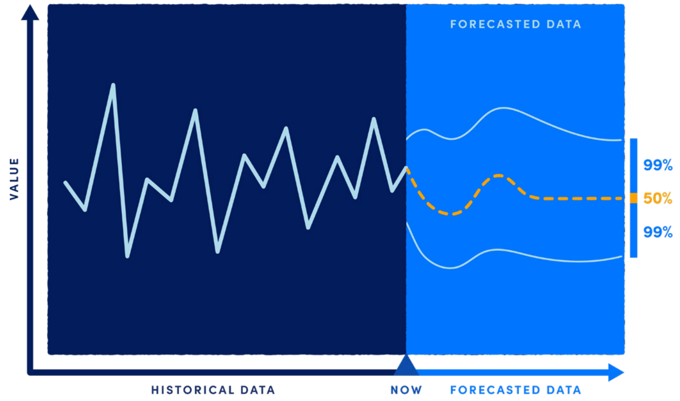

This graph represents a forecasted data range with 99% confidence. Activity represented in the bark blue region is used to understand what future performace looks like and is reflected in the light blue region.

Getting started

AIOps is a transformative initiative that spans your entire IT infrastructure, and profoundly enhances the way your IT and business teams work together. Naturally, you will want to plan carefully as you develop your AIOps initiative.

A good way to get started is to become familiar with ML and AI capabilities right away, even if your AIOps initiative is still on the horizon. You may also want to examine your organization to find data and analytics resources that are available to support your initiative. Set up some initial test cases that will enable you to learn more about what AIOps can do, and apply a step-by-step approach to deployment.

In the next chapter, we'll review some of the steps you can take to move forward on your AIOps journey.

Chapter 3: Implementing your AIOps initiative

AIOps is a great way to empower IT organizations with the visibility and insight they require to spot potential issues, as well as the context and automation they need to prevent incidents. What sets it apart is its ability to bring together data and metrics from multiple systems, then present the most relevant information to enable organizations to prevent problems, instead of reacting to them. But migrating to AIOps is rarely an initiative that is practical with in-house resources.

In this chaper

In Chapter 2, we presented an introduction to some of the basic elements that make up AIOps. In this chapter, we'll look at some of the key considerations involved in building an AIOps system and team. We will discuss some of the most important AIOps processes and platforms, and the importance of automation in addressing the vast amount of data that resides across an organization's technology stack, and surfacing only the most important alerts about anomalies that could potentially cause trouble. We will also explore some of the basic steps and choices involved setting up an AIOps initiative.

In this chapter, you will learn:

- Processes and platforms that support AIOps
- The importance of automation
- How to set up an AIOps initiative

Solve for complexity with AIOps

If we consider how organizations monitored their IT infrastructure a decade ago, it's clear that the world was much simpler. If you were building an enterprise service or application, you had ten or twenty technology options to choose from. When it was time to deploy, if you needed help with set up and support, database administrators (DBAs), engineers, and other domain experts could provide all the assistance you needed.

Today, we have hundreds of solutions to choose from to build applications—and almost no organization has the in-house knowledge to fully support all of them. At the same time, applications have become increasingly distributed, and are using more open source technologies and solutions. Today, it's nearly impossible for a single IT team to attain the level of expertise needed to support every type of technology in use. IT support and domain knowledge that used to be deep and narrow has become wide and shallow. AIOps was developed to help organizations address these challenges.

AIOps processes and platforms

As detailed in previous chapters, AIOps is all about using AI algorithms to make IT Operations tasks easier. By automating analysis, you let the intelligence in these algorithms do the work for you. If you're an IT Operations professional, you will spend less time by automating the analysis that you would normally have to perform manually—and more time focusing on innovation and initiatives that can drive business growth. Let's take a closer look at some of the basic processes required to put effective AIOps into action.

Solving the big data problem

Even organizations with full-time data scientists on staff understand that they are only one part of a much larger process. Garbage in, garbage out is the rule of thumb in data science, so your insights can only be as good as your input. That makes it essential to be sure you are bringing in the most relevant data, from the right flow, the right time period, and for the right event.

In IT Ops, a datapoint is a piece of data that is collected during monitoring. Some of the most common datapoints include data from your infrastructure's devices, cloud resources, instances, websites, services, and groups.

You must not only acquire the data you need, but bring it together into one place, which is often challenging in today's diverse, multi-cloud environments. Before you can proceed with AIOps analysis, you need to know if you've solved your big data problem, bringing together all the relevant data you need, so it can move forward into the next step.

Ensuring meaningful alerts

The second key challenge for AIOps focuses on alerts. Consider whether you trust the alerts that are currently generated by your monitoring products. Are they relevant, providing meaningful context? If they don't, then AIOps capabilities can be applied.

If your alerts are highly distributed, with dozens of tools that are generating a flood of data, you need to take steps to detect the signal through the noise, through dynamic thresholds.

In some cases, you may have difficulty determining exactly what your issue is. You may only know that you are simply seeing too many alerts, creating so much noise that it's difficult to understand what issues are important. One helpful question to ask is how many of your alerts you need to take action on. If 100% of the alerts you are receiving are actually important issues, you already have a strong foundation to build an AIOps layer upon. On the other hand, if only 70% of your alerts are critical, then depending on your organizational size, you will be grappling with a great deal of noise in your environment.

The importance of a single platform solution

The advantage of AIOps capabilities is that it can be applied to provide intelligence to every element of the IT stack. In a successful AIOps approach, all observable data of metrics, logs and applications, should be centralized in one platform. AIOps provides intelligence into a variety of elements, including:

- **Cloud environments,** enabling you to monitor AWS, GCP, VMware, Microsoft Azure, and other cloud services and infrastructure together on one pane of glass
- **On-premises networks,** providing visibility into the health and performance of your networking equipment and other infrastructure
- **Servers,** providing monitoring and alerts for applications, hardware, and OS metrics
- **Websites,** enabling you to maximize website performance and availability and enable synthetic transactions
- **Containers,** providing monitoring for Kubernetes pods, nodes, containers, and other components
- **Business-critical applications,** enabling you to track availability, performance, and advanced metrics
- **Dynamic services,** monitoring today's complex, geographically dispersed and ephemeral resources to fully assess service health
- **Storage,** tracking overall storage system health and individual components
- **Remote monitoring,** ensuring your remote workforce has access to the tools they need to maintain business continuity

For example, in cloud environments, service configuration changes can have an immediate impact on applications. By providing insight into what these service changes are and when they happen, and context into how they relate to application performance, AIOps could help reduce mean time to repair for issues.

By gaining better insight into on-premises networks, an organization could dig deeper into the root cause of issues. If a branch office with a remote desktop environment is experiencing performance issues, AIOps could analyze detailed networks to narrow down the cause, and demonstrate that the issue is related to an overly long CPU queue, rather than issues with other elements of the technology stack. This level of insight might not be possible with conventional monitoring alone.

Getting started with AIOps

If you're ready to consider deploying AIOps across your organization, what are some initial steps you can take to move forward on your journey? Start by evaluating your business processes, your environment, and their associated KPIs.
For example, you may wish to focus on improving MTTR for infrastructure issues. Or you may wish to enable individuals on your IT team to maintain more parts of the environment, in less time. Determine which KPIs are most relevant, and how you will drive the outcomes you are trying to achieve.

Simplicity is key

You shouldn't need to be a data science expert to take advantage of the benefits of AIOps. One way to accelerate your ROI is to choose a solution with dynamic thresholds that are pre-set to a sensible default. Your solution should be relatively simple to turn on and start using right away, using an intuitive interface.

Avoid choosing platforms with AIOps capabilities that behave as though they were designed and built exclusively for data scientists. There is no need to expose every data science parameter as part of its initial setup. An effective platform will enable most to fine-tune the solution to ensure their alerts are relevant and actionable, but the solution should be easy to start using right out of the box.

Monitoring and observability

AIOps is all about acquiring insight into the big picture. Choose a monitoring solution that is capable of monitoring complex infrastructures that have a mix of on-premises, cloud platforms, containers, networks, and storage environments. A solution that embraces observability by being able to monitor metrics, log and application data in one platform. Finally, one that demonstrates the ability to provide more actionable, fruitful results than a product that is only capable of monitoring discrete cloud, network, or container environments.

For example, if you are only monitoring an on-premises infrastructure, your insights will be limited to that environment. If you have an on-premises data center that extends out to the cloud or additional storage environments, you will miss events and issues that span your full environment. AIOps products that offer narrow, siloed views simply can't acquire all the data they need to properly correlate events and automate the presentation of actionable information.

SaaS vs. on-premises solutions

Like most technologies, the monitoring and AIOps domain is evolving all the time, with new technologies and better algorithms emerging constantly. To keep pace with constant enhancements, a software-as-a-service (SaaS) delivery model provides the flexibility and capabilities to fit changing needs.

Compared to an on-premises solution, a SaaS monitoring and AIOps platform is easy to deploy, and is regularly enhanced to take advantage of the very latest technologies. When a new algorithm emerges that can accelerate troubleshooting or enhance process signaling, you can take advantage of it right away, instead of waiting for manual upgrades that can take time and resources to apply to an on-premises product.

As your business needs continue to change and your infrastructure evolves, your monitoring and AIOps platform should also be able to scale with it. Consider how you can save time and realize additional business outcomes by extending your solution across more data types and analyzing more data.

Focus on the outcome

We've seen a great deal of hype around AIOps, and many vendors are developing AIOps solutions simply to show that they have one on the market. But to be effective, an AIOps solution should be built to focus on providing an early warning system and failure prevention system. It must first determine what is important in an environment, then enable an organization to add more context to the data it monitors, and make it more relevant and actionable. The faster you can detect and resolve issues, the more proactive and agile your organization will be.

Simply put, AIOps lets you shift to preventing issues instead of simply responding to them. That's critical in a dynamic, competitive environment where agility and rapid go-to-market is key.

Part 2

How to use AIOps for monitoring

Understanding the core elements of an early warning system

Introduction

An early warning system is the foundation of AIOps for monitoring. In this section, we'll explore the fundamental components that form an early warning system: dynamic thresholds, root cause analysis, and forecasting. Each chapter is dedicated to each functionality, how it is built, and the business benefits it offers. Together, the combination of these three elements allow you to surface only the anomalies and alerts that are most relevant and bring into focus the ones that require action.

With an early warning system in place, your organization will be ready to unlock the true power of AIOps for monitoring: automated failure prevention.

Chapter 4:
Early warning system overview & benefits

An early warning system offers significant benefits for cost optimization, risk mitigation, operational efficiency, and revenue growth. As you consider adopting an AIOps approach in your IT Organization, it's important to understand the foundational components - why they are necessary, how they work, and the business benefits your organization stands to gain when correctly implemented.

In this chapter

We will take a deeper dive into the different components of an AIOps early warning system. We will examine how dynamic thresholds, algorithmic root cause analysis, predictive forecasting and automated anomaly detection of unexpected performance and change events give IT professionals the head start they need to take fast, informed action based on the context of the situation unfolding.

In this chapter, you will learn:

- The importance of an early warning system
- How the EWS can prevent downtime
- How to prevent alert fatigue

Early warning system benefits for complex environments

Today's modern hybrid infrastructures are increasingly diverse, distributed, multi-cloud environments. They often have a mix of some resources in the cloud, with other resources running on-premises in more traditional data centers. To allow for faster deployment and change, immediate scalability based on user needs, and to standardize application delivery across cloud and on-prem environments, organizations are also increasingly running their applications in containers, more than half of which are doing so in hybrid mode.[11]

These mixed, modern environments are great for enterprises, because they enable them to fine-tune their environments and workloads for the best cost, security, scalability, deployment speed, and other requirements. But the catch is that these diverse environments are also much more complex to monitor.

Historically, most monitoring products have specialized in monitoring one element of the IT stack well, such as servers, databases, logs, or cloud resources.However, these products were not intended to monitor the full-stack in a unified, contextual way. Any change to the organization, such as the acquisition of a new cloud service, a new on-prem infrastructure technology, or a new strategic application, would usually prompt investment in an entirely new monitoring product or "tool." IT Ops were then required to learn new skills or hire people who could use these products. It's common to see enterprises using multiple monitoring products, and even dedicating full-time staff to configuring and managing the tsunami of data produced.

As infrastructures continue to expand, these additive approaches to IT operations can't continue to scale. When IT teams are spending most of their time reacting to problems and minimizing downtime, they are less able to focus on digital transformation and innovation. Digital businesses have begun to realize that monitoring is not only about alerting on a particular aspect of the IT stack, but about changing data into actionable insight, via AIOps capabilities. That's leading to an increasing strategic focus on unified observability, and the ability to be able to collect and make sense of metric, application and log data in one platform.

[11] Kubernetes and Container Security and Adoption Trends [StackRox], 2020

Surfacing hidden Signals to power rapid action

An effective AIOps system will detect and provide early warnings about the signs and symptoms that precede issues, such as patterns or anomalies in performance or log data, and warn users about them in advance. These early warnings keep IT better informed, by employing automatic anomaly detection via ML or powerful algorithms that detect issues and their root cause, before any widespread business impact. It's the first step to moving from a proactive to a predictive monitoring mindset.

At the heart of the early warning system are AI and Machine Learning (ML) algorithms that support anomaly detection in any data set, such as IT infrastructure metric and log data; dynamic thresholds; root cause analysis; and automatic correlation. Working together, these features sift through massive amounts of monitored data (e.g. metric and log events) and surface the most important information, then make it more actionable by adding context.

They can also support advanced log analysis, where contextual and actionable log data is correlated with metrics and alerts to provide deep insight that goes beyond basic notifications—to help you understand why issues are occurring. An AIOps early warning system should help prevent major business-impacting incidents, by processing the signal through a rule-based action engine, tied into a robust automation framework. The result is a more informed, proactive IT operations team, and shorter MTTR.

Preventing downtime

One of the most compelling benefits of an early warning system is its ability to minimize downtime. According to a recent study[12], companies with frequent outages and brownouts experience 16 times higher costs than organizations that have fewer instances of downtime. These organizations also require twice the number of team members to troubleshoot problems—and spend twice the amount of time doing so. The impact is not limited to IT. When you consider the cost to an organization in terms of lost revenue, a poor customer experience, or damaged reputation, the consequences of unplanned downtime quickly multiply.

Downtime not only creates immediate expenses but reduces momentum for strategic initiatives and driving business growth. When IT teams are spending most of their time tracking down root causes and putting out fires, they do not have the time or resources to switch to a proactive model where the focus is on avoiding outages instead of minimizing their impact. An AIOps early warning system will provide IT operations with the information they need to make the switch to proactively preventing problems, instead of reacting to them.

An AIOps early warning system solves the problem of traditional monitoring approaches that focus on static alerting and analysis configurations such as static thresholds. As environments have become dynamic and distributed, these types of traditional approaches weren't practical, because they required human intelligence and significant time and effort—something that today's digital businesses can't afford.

Minimizing alert fatigue

Alerting is an essential aspect of preventing downtime, but it can also be one of the most frustrating, time-consuming parts of an IT professional's job. Gaining better visibility into infrastructure monitoring can be a double-edged sword. It's one thing to acquire insight into everything that's going on across all your environments, but if your IT teams are in situations where they are constantly bombarded with multiple alerts due to poorly-tuned alert thresholds, alert fatigue can rapidly set in. According to a recent survey, 47% of organizations experience more than 50,000 alerts per month.[13] When the sheer volume of alerts—relevant or not—begins to overwhelm IT professionals, people begin to get burned out. Alerts may be ignored, responses slow, and incident management gradually gets worse.

With its ability to surface the most relevant alerts from a broad array of data, an early warning system classifies alerts and workloads more effectively. Using dynamic thresholds, it detects the normal performance range for technical and business metrics, and generates alerts based on anomalies. It then alerts IT teams based on historical performance and advanced algorithms. It helps businesses avoid alert fatigue, save time, and surface anomalies sooner.

[12] 2019 IT Outage Impact Study [LogicMonitor], 2019

[13] Paige, Viki. Reducing Alert Fatigue: How Your Automation COE Can Help [Broadcom], Jun. 2020

Opening up visibility, context, and understanding

Ultimately, your infrastructure is intended to power better business outcomes by supporting innovation, accelerating time to market, and delivering a superior customer experience. But our ability to realize these outcomes is directly tied to the health of your environment.

An effective AIOps solution should provide full insight across the technology stack, from the cloud to traditional on-premises systems to containers and the applications running within them. This comprehensive visibility enables an AIOps early warning system to identify and help prevent issues across the entire complex, distributed modern IT infrastructure to maximize availability of the resources that drive the business.

Preventing widespread business problems instead of reacting to them

An AIOps early warning system is the game-changer that lets you break the cycle of constantly chasing down issues after they happen. By preventing issues from becoming widespread in their business impact, it makes the flood of data your infrastructure environment is producing more manageable and actionable, and frees your team to focus on the big picture. In our next chapter, we will explore dynamic thresholds, a key element that make up the AIOps early warning system, and illustrate how they benefit IT professionals at every level of your organization.

Chapter 5: Dynamic thresholds

In our previous chapter, we presented an overview of the early warning system and its ability to rapidly surface anomalies to bring into focus the most urgent alerts. Now, we will explore the dynamic thresholds that make up a key component of the early warning system.

In this chapter

We will provide an overview of the algorithms that enable you to discover and act on what's truly important, which is extremely challenging even in today's complex, dispersed multi-cloud environments. We will also discuss how gaining this proactive insight enables IT leaders to act more strategically, and help their IT teams to execute more effectively.

In this chapter, you will learn:

- How AI/ML-based algorithms detect anomalies
- How dynamic thresholds set the stage for automatic remediation
- Why dynamic thresholds are a requirement for proactive it

Discover the exceptions with dynamic thresholds

An early warning system is about much more than simply monitoring complex, distributed, and ephemeral environments and alerting in context. To be really effective, you first need to determine which data points are most important, separate them from those less relevant, and notify your IT team about them, before they impact performance or availability. That process starts with improving focus utilizing dynamic thresholds.

Dynamic thresholds are based on AI/ML-based algorithms focusing on anomaly detection based on rate of change and seasonality, along with algorithms to contextualize issues. These algorithms automatically detect the normal performance range for any metric—whether it's a technical or business metric—and accurately alert based on values outside of this range that are considered anomalies.

Because dynamic thresholds (and their resulting alerts) are automatically and algorithmically determined based on the history of a datapoint, they are well suited for datapoints where static thresholds are hard to identify, such as monitoring the number of connections, latency, and other criteria. Dynamic thresholds are also useful in situations where acceptable datapoint values aren't necessarily uniform across an environment.

As cloud deployments increase and environments become more volatile, dynamic thresholds also provide a more effective way to discover exceptions in these ephemeral environments. While static thresholds are too difficult and time-consuming for people to manage manually, dynamic thresholds provide an opportunity for AI/ML-based algorithms to demonstrate their advantages in fast-changing environments.

Identifying patterns and pattern types

When most people hear about dynamic thresholds in the context of monitoring, they think of the ability to simply trigger alerts based on a non-static threshold. Although this capability is incredibly useful, it's only half the picture.

An effective solution should enable you to do much more. It should not only trigger alerts by identifying issues that wouldn't be caught by traditional static thresholds. It should also help you focus, by eliminating excess alerts caused by thresholds that haven't been tuned well, surfacing the important issues. This lets you use dynamic thresholds as a filter to reduce alerts to only what is immediately relevant.

Simply put, dynamic thresholds represent the bounds of an expected performance range for a particular datapoint. Even the bounds themselves are subject to change—for example, a CPU running at 85% capacity might be abnormal for certain workloads, but normal for others. Unlike static thresholds that are assigned manually, dynamic thresholds are calculated by anomaly detection algorithms, and continuously trained by a datapoint's recent historical values.

When you enable dynamic thresholds, powerful algorithms calculate the thresholds to be set, and continually adapt to environmental changes, both from a technical and business perspective. In other words, alerts are generated when anomalous performance is detected.These dynamic threshold algorithms can detect a variety of different types of data patterns, including:

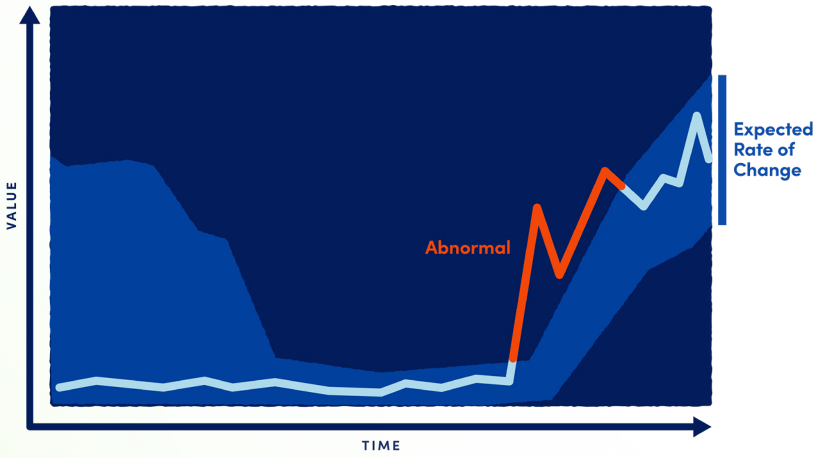

In this diagram, the light blue shaded area indicates the expected rate of change for a particular metric. This means that an alert will not be generated when the metric is in this range. However, the red line indicates that the monitored metric has moved outside of the expected rate of change and should be investigated.

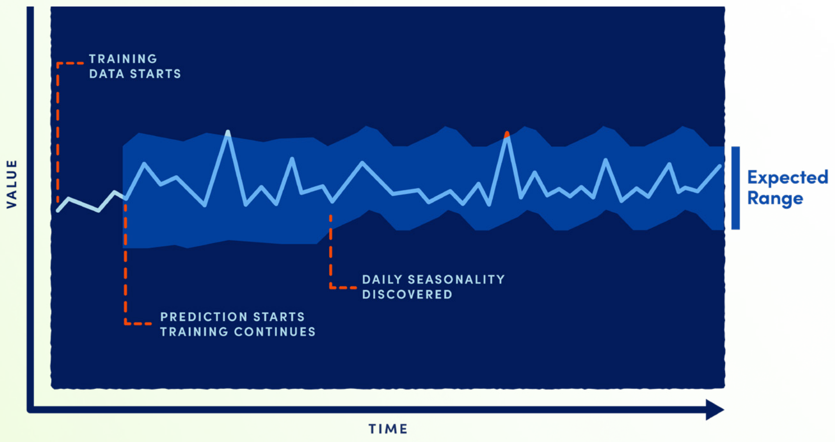

In this diagram, a daily seasonality pattern has been discovered and learnt by the Dynamic Threshold algorithm. Combining this with rate of change calculations means that alerts are only relevant if the metric deviates from the expected range (light blue area).

Anomaly detection

To discover anomalies, historical performance is used to generate an expected range for resources. This range is used as the basis to highlight and alert on activity that breaches this range and is considered anomalous. With anomaly detection, activity that may have previously triggered an alert is now silenced based on historical performance.

Rate of change

Anomaly detection is not informed by metric value alone, but also enables dynamic thresholds to alert on anomalies in metric value rate of change. For example, monitoring could produce a significant increase in a metric that has not yet crossed a static threshold, such as a disk that is rapidly filling up. Providing insight into a metric's rate of change enables IT teams to catch issues sooner, before they negatively impact the business.

Seasonality

Dynamic thresholds should map seasonality for your business. Algorithms ensure that your AIOps solution identifies patterns in resource performance, and alerts on anomalies in these patterns. For example, you might have a VM in place that normally backs up on a daily basis but has failed to back up on a particular day. An early warning system could alert you to the issue right away, well before it becomes a more serious issue.

A well-designed AIOps solution should present an auto-generated range that is widely applicable with sensible defaults, while enabling you to customize it if needed with advanced configuration options.

Benefits that span the IT team

Under an inefficient, poorly tuned alert system, issues can quickly escalate from manageable problems to fire drills. Solid dynamic threshold features should help avoid these issues. As part of an AIOps approach, you may utilize pre-set best practices thresholds that will filter out noisy alerts, and raise issues for a typical production environment. This positions IT professionals at all levels to achieve significant gains in efficiency and performance.

Operational efficiency

Customizing thresholds is often time-consuming and difficult for large environments—and becomes nearly impossible when you begin to add cloud and ephemeral environments to the picture. Dynamic thresholds help you maximize IT efficiency by ensuring that alerts are only sent for anomalies that matter from a user (employee or customer) and business perspective, even in the most complex infrastructure environments. Accurate, early alerts also help build trust that enables you to move towards automatic remediation—the failure prevention system.

Improving IT efficiency eliminates the need for manual management of monitoring. It lets you free up your valuable resources for other tasks, while enabling you to realize a better monitoring ROI.

Issue prevention

Dynamic thresholds enable your teams to understand the correct level of expected performance—and immediately spot where it deviates from normal and requires attention.

At the same time, they help you ensure that your teams aren't receiving irrelevant alerts for optimized machines that are regularly highly utilized.

In situations where it is important to determine if a returned metric is anomalous, dynamic thresholds have tremendous value. Not only will they trigger more accurate alerts, but in many cases issues are caught sooner, before widespread business impact. In addition, administrative effort is reduced considerably because dynamic thresholds require neither manual upfront configuration nor ongoing tuning.

Benefits for IT teams

IT teams such as system engineers, analysts, and infrastructure architects, are responsible for increasingly complex, dispersed infrastructures. If you're part of an IT team working on the front lines, it is difficult to gain visibility and control when your infrastructure and apps extend across on-prem, private, and public cloud environments. For front-line warriors, alert fatigue is a primary concern. Your team is likely overloaded with alerts when static thresholds don't exist or are not tuned properly. That can easily lead to missed or ignored alerts. When an issue does arise, the vast volume of alerts can make it extra time-consuming to fully understand what's happening and take the right steps to address it.

Static thresholds can help sort out some potential issues, but they are cumbersome to set up and maintain. Accurately, specifically tuning static thresholds for a large environment requires a great deal of administrative overhead. The challenges only grow in environments where resources are utilized differently. Some resources may be highly utilized intentionally, while others are not.

Dynamic thresholds help you save time customizing static thresholds, because they employ a dynamic expected range to determine whether alerts should be routed. Alerts within the range are still displayed within the AIOps platform, but users are not notified about them. As part of an early warning system, dynamic thresholds will automatically identify the normal resource consumption, even for an intentionally highly utilized resource, and only notify users of consumption that is not normal. This lets them focus on addressing actual issues, instead of investigating minor variances.

Benefits for IT leaders

If you are an IT leader such as a manager or director, it is likely you face many of the same challenges as individuals on IT teams. You are not only responsible for complex, dispersed infrastructures, but are also expected to collaborate closely with senior executives to drive digital transformation and other strategic initiatives.

Ultimately, it's up to you to ensure maximum performance and consistent availability for your organization's infrastructures. To do it, you need a large, highly skilled team to set up and manage alert tuning and static thresholds.

Because long term, strategic thinking is such a key part of their role, IT leaders like you are also up against another type of challenge: planning and gaining visibility into future performance of their infrastructures.

An AIOps platform enabled by dynamic thresholds is a powerful way to de-risk operations as you evolve IT. By enabling your teams to expedite MTTR for new issues, it gives you the breathing room you require to adopt next gen technologies, drive innovation, and refine better, long-term strategies.

Chapter 6:
Root cause analysis

In our previous chapter, we discussed how dynamic thresholds apply AI/ML algorithms to spot technical and business metrics that are outside normal performance. These dynamic thresholds enable organizations to surface the most important, relevant alerts, and detect anomalies within a variety of different types of data patterns.

In this chapter

Now we will take a closer look at how AIOps can dig deeper to provide insights into the events that are most important. The root cause analysis (RCA) feature in an AIOps early warning system delivers the clarity and the control over alert notifications that are required to make it happen.

RCA identifies the root cause when an issue occurs, enabling IT operations engineers to focus on solving the issue quickly instead of wasting time searching for it. When you combine RCA with an AIOps platform's ability to monitor most anything, including containers, cloud environments, network, storage, and more, it becomes possible to reduce downtime even for highly complex hybrid infrastructures.

In this chapter, you will learn:

- Why topology mapping is key and how it works
- 7 critical steps to identifying the root cause
- Benefits for IT teams and leaders

Identifying the true source of an issue

Root cause analysis is a fundamental pillar of the AIOps early warning system because of its ability to not only track down the source of issues, but also filter out alerts that aren't related to that source. It leverages the relationships among your monitored resources to determine the root cause of an incident that is impacting dependent resources. These relationships are best discovered with topology mapping.

For alerting operations, root cause analysis highlights the originating cause of an incident. It also gives you the option to suppress notification routing for alerts that are dependent on the originating alert. This dramatically reduces alert noise for events in which a parent resource has gone down or becomes unreachable, and causes its dependent resources to go into alert as well.

Identifying the root cause is the first step in enabling a more proactive approach to your infrastructure, through a failure prevention system. An AIOps failure prevention system should also give you the ability to automate actions that remediate the root cause issue. This closes the loop from intelligently identifying and predicting issues to automatically fixing and preventing them. The result is not only reduced downtime for your infrastructure, but more free time that your IT operations teams can spend on innovating and transforming your business.

Using topology mapping to find the root cause

Performing root cause analysis requires an understanding of the relationships between monitored resources. Ideally, these relationships should be automatically discovered via topology mapping. Topology mapping is the visual representation of relationships among elements in a communications network. These maps can represent the physical location of network components, generally referred to as layer 1 mapping, or they can represent the logical relationships among elements, referred to as layer 2 mapping.

With AIOps, topology mapping capabilities should focus on layer 2 mapping. For example, LogicMonitor, an intelligent unified observability platform, utilizes the Link Layer Discovery Protocol (LLDP) as well as Cisco's proprietary version of the protocol known as Cisco Discovery Protocol (CDP) to dynamically generate network topology maps that show how data flows among the different resources like switches, hosts, firewalls, routers, and other network components in your environment.

With an additive approach to generating topology maps, only the relationships that are most relevant for the current task are shown, but you can expand the view outward.

One of the most powerful benefits that topology mapping provides is context. For RCA, it should give you the ability to quickly determine the root cause of an incident that is impacting dependent resources. It should also provide the option to suppress alert notification routing for alerts that are determined to be dependent on the originating alert.

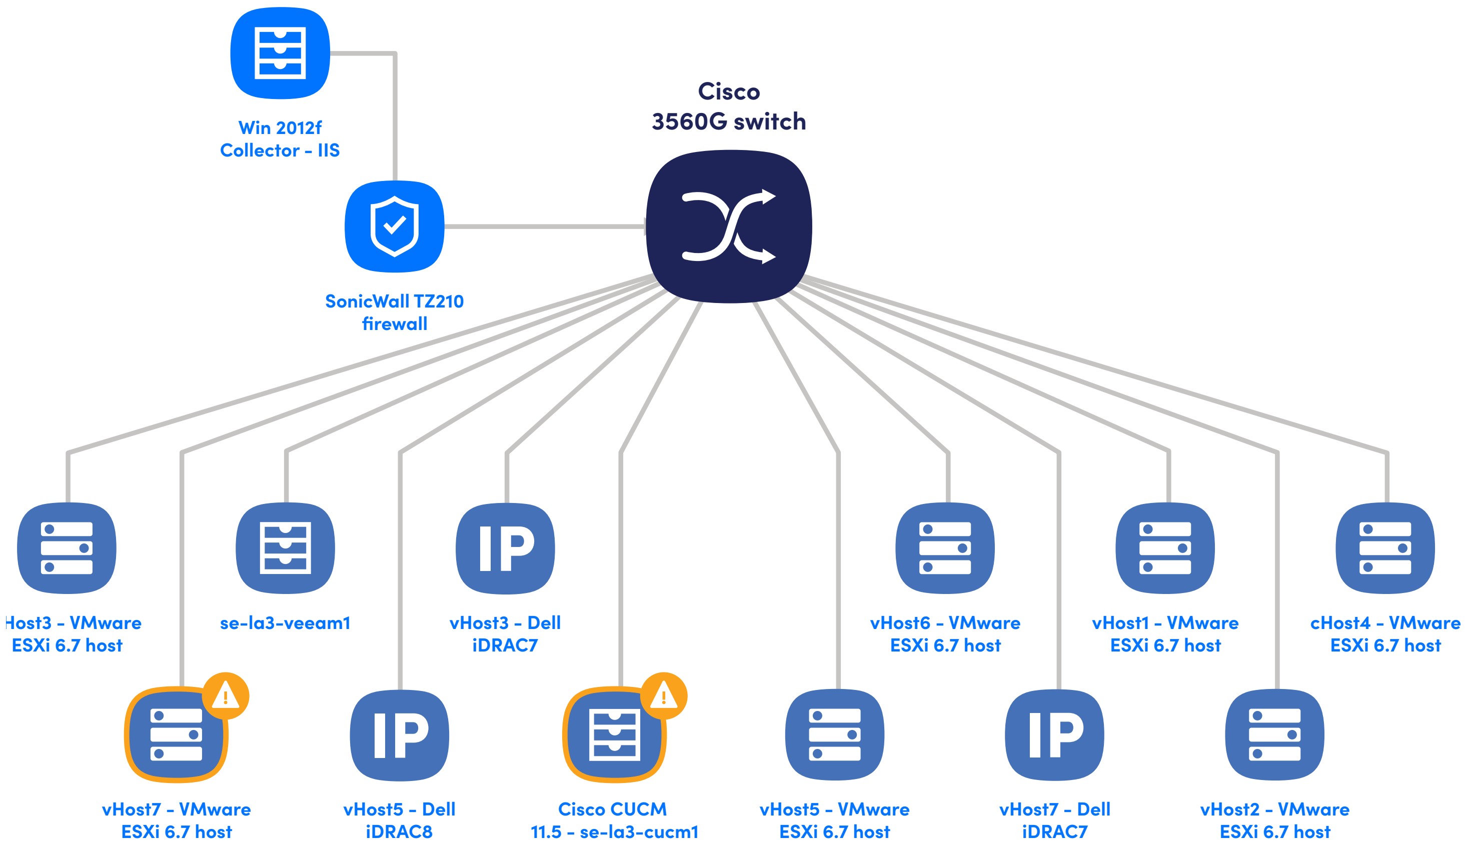
Win 2012f Collector - IIS
Cisco 3560G switch
SonicWall TZ210 firewall
Host3 - VMware ESXi 6.7 host
se-la3-veeam1
IP
vHost3 - Dell iDRAC7
vHost6 - VMware ESXi 6.7 host
vHost1 - VMware ESXi 6.7 host
cHost4 - VMware ESXi 6.7 host
vHost7 - VMware ESXi 6.7 host
IP
vHost5 - Dell iDRAC8
Cisco CUCM 11.5 - se-la3-cucm1
vHost5 - VMware ESXi 6.7 host
IP
vHost7 - Dell iDRAC7
vHost2 - VMware ESXi 6.7 host

Topology mapping is the visual representation of relationships among elements in your communications network.

How AIOps-based root cause analysis works

During an alert storm, your IT ops teams may encounter multiple alerts relating to the same originating incident. Each time a metric threshold is exceeded, a notification is sent, and a trickle of notifications can quickly become a wave. Before long, you're up against a flood of notifications for every resource that's affected by the incident—without having a clear idea of which resources are its true cause. Here's how enabling root cause analysis addresses this issue:

1. Identify unreachable alerts for resources in a dependency chain

Root cause analysis is based on topology relationships. If a resource that is part of an identified dependency chain goes down or becomes unreachable, its alerts are flagged for root cause analysis.

2. Delay routing of alert notifications

When a resource in the dependency chain goes down or becomes unreachable, this first "reachability" alert triggers all resources in the chain to enter a delayed notification state. When enabled, this state stops alert notifications from being immediately routed, which buys time for the incident to fully manifest. It also frees up time to enable the root cause analysis algorithm to determine the originating causes, and the dependent causes.

3. Add dependency role metadata to alerts

In this step, any resource in the dependency chain with a reachability alert is identified as a parent node or suppressing node to its dependent child or suppressed nodes. This process adds metadata to the alert, and specifies whether its role is originating or dependent. This role provides the data that AIOps to suppress dependent alert notifications. Root cause analysis can also put the right log event in the context of an alert, automatically, to further enrich responses.

4. Suppress routing of alert notifications

For this optional function, alerts identified as dependent are not routed. This enables RCA to reduce alert noise reporting only those alerts that identify root cause. You may still view dependent alerts on your AIOps platform interface, but your team won't receive notifications.

5. Clear alerts across dependency chain

When the originating reachability alerts begin to clear, all resources in the dependency chain are once again placed into a delayed notification state. This allows time for the entire incident to clear. After a few minutes, any remaining alerts will then be routed for notification. If some resources are still unreachable, RCA initiates a new root cause analysis incident, and the process repeats.

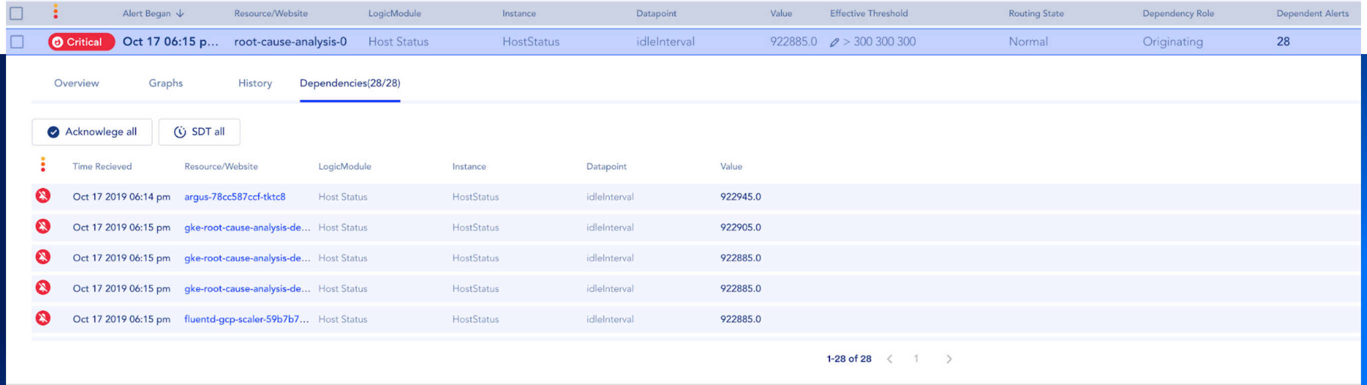

□		Alert Began ↓	Resource/Website	LogicModule	Instance	Datapoint	Value	Effective Threshold	Routing State	Dependency Role	Dependent Alerts
□	Critical	Oct 17 06:15 p…	root-cause-analysis-0	Host Status	HostStatus	idleInterval	922885.0	> 300 300 300	Normal	Originating	28

	Time Recieved	Resource/Website	LogicModule	Instance	Datapoint	Value
	Oct 17 2019 06:14 pm	argus-78cc587ccf-tktc8	Host Status	HostStatus	idleInterval	922945.0
	Oct 17 2019 06:15 pm	gke-root-cause-analysis-de…	Host Status	HostStatus	idleInterval	922905.0
	Oct 17 2019 06:15 pm	gke-root-cause-analysis-de…	Host Status	HostStatus	idleInterval	922885.0
	Oct 17 2019 06:15 pm	gke-root-cause-analysis-de…	Host Status	HostStatus	idleInterval	922885.0
	Oct 17 2019 06:15 pm	fluentd-gcp-scaler-59b7b7…	Host Status	HostStatus	idleInterval	922885.0

Configuring root cause analysis

Every set of RCA configurations you create should be associated with one or more entry points. An entry point is the resource where the dependency chain begins—for example, the highest-level resource in the resulting dependency chain hierarchy.
All resources connected to the entry-point resource become part of the dependency chain. That means they're subject to root cause analysis if any device upstream or downstream in the dependency chain becomes unreachable.

You can configure different settings for different entry points, for flexibility in setting up your root cause analysis system. For example, a managed service provider (MSP) may have some clients that permit a notification delay, but others that don't, due to strict SLAs. Or an enterprise may want to route dependent alerts for certain resources, but not others.

Selecting an entry point resource establishes a dependency hierarchy. In this hierarchy, every connected resource is dependent on the entry point, as well as on any other connected resource that is closer than it is to the entry point. This means that the triggering of root cause analysis does not happen only when an entry point becomes unreachable and goes into alert. Any node in the dependency chain that is unreachable and goes into alert will trigger root cause analysis.

A robust AIOps solution can also support contextual root cause analysis, to put different elements of observable, monitored data into context. Correlating IT infrastructure metric data with the right log event data can dramatically enhance the capabilities of root cause analysis, especially when it comes to change-based anomalies.

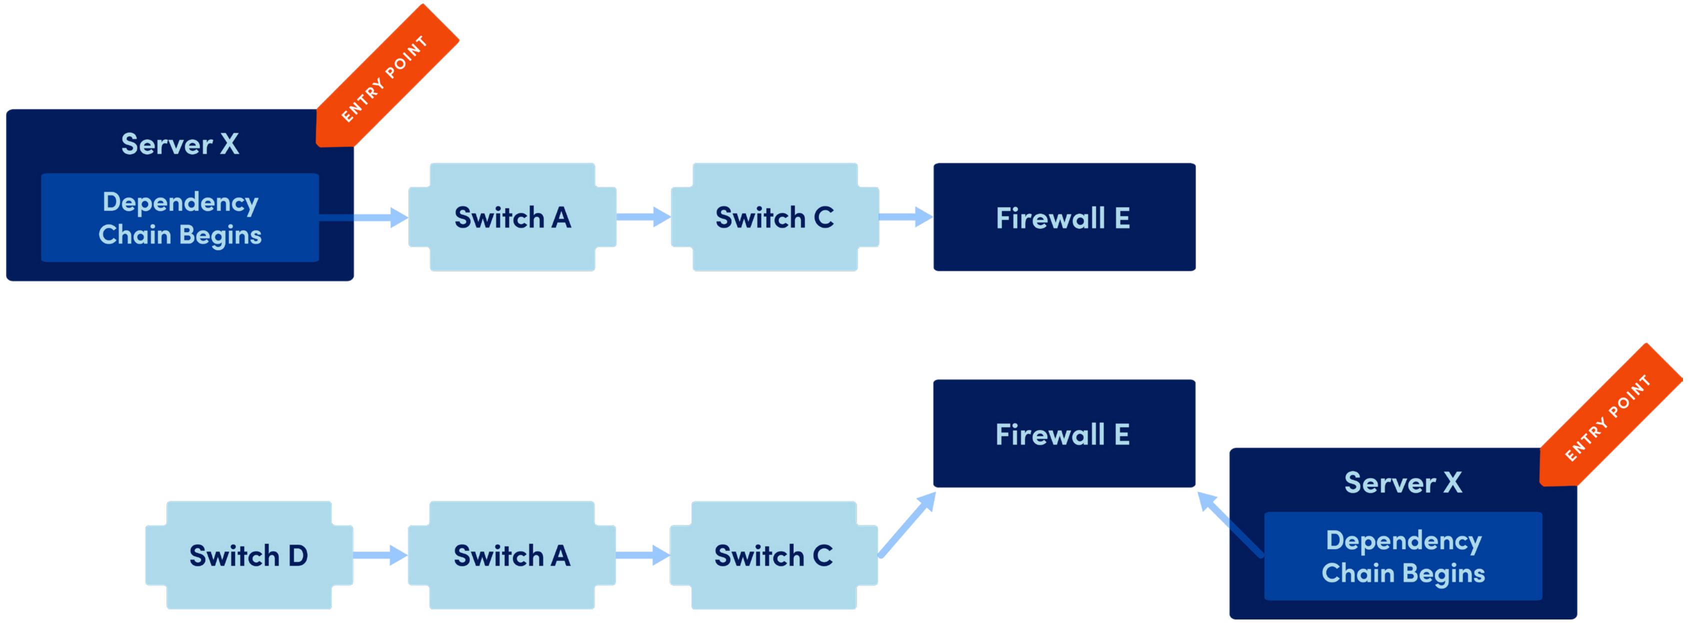

In the above diagram, you can see the logical dependency chain starts with Server X. This is the entry point in which RCA configuration should be applied, as it's the start of the dependency relationship. An issue with Server X will cascade to all dependent resources.

Benefits for IT teams

For IT teams, automated root cause analysis is an essential capability for saving time and separating signal from noise in alert notifications. Every day, analysts, system engineers, and architects are grappling with complex on-prem, private, and public cloud environments. They're doing all they can just to keep up, and any relief they can have from putting out fires means more time they can spend focusing on their customers or working with IT leaders and business stakeholders on more strategic projects. As part of an AIOps early warning system, root cause analysis can help you accelerate mean time to resolution (MTTR), and minimize alert fatigue.

Speeding mean time to resolution

IT teams know that understanding the context of an IT incident can provide the insight needed to greatly reduce the MTTR—and enhance the ability to determine the root cause. In some situations, this context might be the downstream dependencies after a high availability pair of firewalls goes offline.

In other scenarios, the context might be the datastore in contention from multiple VMs. No matter what the issue, context offers critical information necessary to understand the relevant scope when an incident occurs.

The topology mapping functionality that's part of root cause analysis gives IT teams the full context needed to troubleshoot events or incidents that may have occurred, and do it more quickly. Log analysis can also provide valuable insights into how issues are occurring, to speed issue resolution. That means more uptime for critical business systems and processes, and more time for your IT team to focus on more valuable tasks.

Minimizing alert fatigue

Alert fatigue is a very real issue for IT operations teams. If your IT teams are spending all their time chasing down incidents, employee satisfaction drops, frustration sets in, and communication in the team can falter, making it less effective. Alert fatigue can also result in a tremendous opportunity cost. When IT is exhausted by the alerts process, they have that much less time to come up with new ideas, offer a better customer experience, or take proactive steps to prevent future issues.

Root cause analysis does the heavy lifting for you in surfacing actionable alerts within an early warning system. By reporting only on the primary cause of issues, and screening out irrelevant data according to the criteria you choose, it rapidly surfaces any anomalies, and frees your staff from time-consuming, manual work of sorting out what's important and what is not.

Putting the intelligence in AIOps

Root cause analysis brings a high level of insight and context to the AIOps early warning system, to enable IT professionals to see better and take action with more confidence. Our next chapter will provide a look ahead. We'll talk about how forecasting can help you get out in front of issues and predict future trends for your monitored infrastructure, using past performance as the basis.

Chapter 7: Forecasting

In our previous chapter, we focused on the role of automated root cause analysis, and its ability to not only zero in on the source of issues, but also shield IT professionals from the excessive alert notifications that can slow response and contribute to alert fatigue.

AIOps provides strong advantages in rapid troubleshooting, but its capabilities go far beyond tracking and monitoring. It is also an effective way to predict future trends for your monitored infrastructure, using past performance as the basis.

In this chapter

This chapter will discuss how AIOps supports forecasting to enable you to move from proactive to predictive monitoring, mitigating issues before they occur.

In this chapter, you will learn:

- How forecasting applies AI/ML algorithms
- How to use forecasting graphs more effectively
- How forecasting is used to streamline the planning process

Demand and capacity planning data forecasting

The wealth of data that's in your monitored infrastructure not only enables you to respond faster to issues as they're happening; it also enables you to predict where you're headed, especially from demand and capacity perspectives. When forecasting data, effective AIOps capabilities should first identify and remove anomalies and missing data from the set of data you choose for forecasting. Then it applies a capacity trending algorithm to this sample to find a model of best fit for the collected data which then calculates future data based on these model parameters.

Forecasting is an AIOps feature that is especially helpful for discovering and mitigating issues before they impact services. When you pair it with anomaly detection, it helps you to understand what requires immediate attention, or will require attention in the near future.

Consider what would happen if you receive a warning alert indicating that disk usage is at 85%. Your top question will probably be, "how much time until the disk hits 95% usage?" You might have plenty of time to act before disk usage reaches this critical threshold, or it could happen in just a day or two. Forecasting can help you analyze the past rate of rise so you can predict the future rate.

Forecasting is also helpful for budget planning and resource management. For infrastructure components that have lifetimes (such as solid state drives) or capacity (such as disk space) associated with them, forecasting based on the predicted health and performance of your monitored devices can give you insight into the timeframe and magnitude of recurring events. It can also help you anticipate and plan for upcoming expenses.

For example, as IT rapidly embraces cloud computing, AWS platform has become a major portion of many organizations' budgets. Using alert forecasting on an AWS Billing datasource lets you project what your monthly or quarterly spend will be. This is an invaluable tool for everyone involved in drafting your IT budget, including CFOs, CIOs, and Directors.

Viewing and understanding forecast graphs

Just as good AIOps capabilities allow you to visualize anomalies, you should also be able to view data forecasting in an intuitive graphic format. In this example below, the forecast graph uses background color to differentiate between previously-collected data and forecasted data. AIOps forecast graphs provide several different capabilities that you can use to derive additional context from the data presented. For example, they can enable you to control:

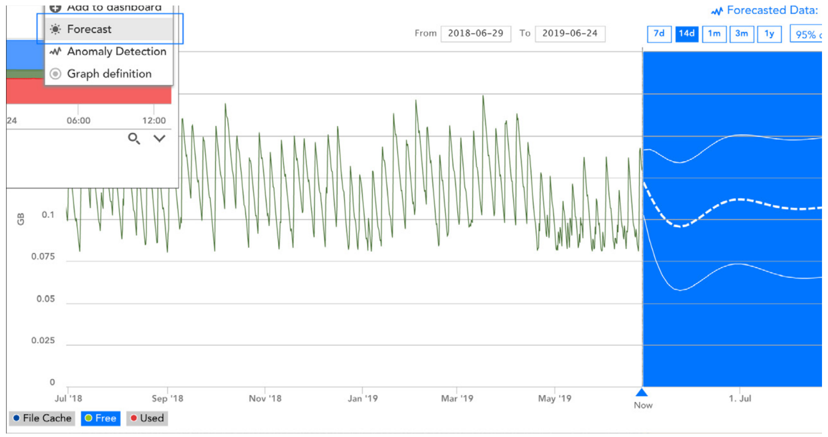

From 2018-06-29 To 2019-06-24

Time range for collected data

You may wish to select a time range for the collected data, such as training data, that will be used as the basis for the forecast. Your solution should enable you to, and one year. You should be able to apply data from the time range you choose to calculate the forecasted data.

Time range for forecasted data

You should also be able to select the duration of the forecast, choosing to forecast out seven days, 14 days, one month, three months, or even a year forward.

Forecast method

An effective AIOps solution will also offer different approaches that you can employ to calculate and display your forecast, such as:

- **95% confidence forecast.** This method projects lines that represent forecast values as well as upper and lower confidence bounds. These confidence bounds indicate that AIOps is 95% confident that future datapoint values will fall within this range.
- **Line of best fit forecast.** This method draws a single straight line that best fits forecasted data. This is generally a more useful option if your data points are highly variable.

Streamlining the planning process

With its ability to help you spot issues before they trigger alerts, forecasting is an essential tool to minimizing downtime that could impact your business, your customers, and your reputation. It can also provide insight into future network health issues, to help you plan more efficiently, and anticipate future time and expenses.

Part 3

The future of AIOps

AIOps use cases for today and tomorrow

Introduction

AIOps enables IT organizations to work more proactively and perform more efficiently through automation and observability.

This section will illustrate how AIOps is being used across different industries and environments today, and point the way forward toward future applications. We will discuss new capabilities on the horizon, such as the failure prevention system, and explore some of the trends that are shaping AIOps development.

Chapter 8: Putting AIOps use cases into action

It’s clear that AIOps is becoming essential for today’s increasingly diverse, cloud-based infrastructures and DevOps environments.

AIOps enables IT to address its increasingly strategic responsibilities and challenges, providing the insights and automation needed to discover and resolve issues faster, and make better, more proactive solutions. The basic components of an AIOps early warning system include anomaly detection, dynamic thresholds, root cause analysis, and forecasting. These elements work together to help IT teams and their leaders deliver superior service quality and availability, collaborate and innovate better, and keep their technology aligned to the business outcomes they want most.

In this chapter

In this chapter, we’ll build on the fundamentals of AIOps—the early warning system components of dynamic thresholds, root cause analysis, and forecasting. We will discuss:

- Real world use cases for AIOps
- Benefits for IT operations
- How AIOps supports a devops environment

Empowering IT operations with the big picture

If you’re an IT operations professional, you know that when an incident occurs and your network or service is impacted, it’s typically the result of a chain of events. From a security breach to a complete system outage, a problem with one service can easily spread to impact another service. Before long, you’re facing a problem that’s compromising overall availability or performance, which ultimately damages your customer’s experience.

When a serious incident hits, your team’s immediate response is to focus on identifying the root cause and restoring service. Because the chain of events for outages often involves a combination of technical and process issues, it can be hard to identify the root cause and understand why the issue occurred in the first place. AIOps connects the dots and develops the context needed to understand the root cause of issues faster, so you can spend less time troubleshooting and more time driving innovation and meeting customer needs.

For example, suppose a site reliability engineer (SRE) on an IT Ops team receives an alert about an application, indicating that requests are failing. Normally, the SRE would have to start manually checking logs or monitoring dashboards to chase down the causes of the issue. An AIOps system could enable the SRE to sidestep some of this diagnostic work, by automatically detecting and surfacing anomalous log events showing that requests to a third-party database service are timing out.

The SRE may find that a recent application configuration change is still referencing, incorrectly, the previously published IP address ranges for the third-party service. By updating the IP addresses directly, the SRE solves the issue in minutes. AIOps has helped the engineer spend far less time troubleshooting, achieve faster MTTR, and work more efficiently, so she can focus on more important responsibilities.

Enabling managed service providers to enhance the customer experience

Delivering a superior customer experience is a key differentiator for today's competitive managed service provider (MSP) market. Because they are required to monitor hardware and performance across multiple, diverse, customer environments and services, AIOps can deliver the visibility and insights that today's MSPs require to support customers more effectively.

For example, a growing MSP in the Netherlands provides deployment and management of ICT infrastructure including networking and Wi-Fi, cloud services, IP telephony, hardware, and software. However, its existing management tools lacked performance monitoring, and provided only limited ability to monitor the latest generation of server hardware, and no options for configuration management. Deploying an IT infrastructure monitoring and observability platform, AIOps capabilities, enabled the MSP to gain visibility into its entire environment and take advantage of data collection sets, pre-set activity thresholds, and automated alerts.

With so much more insight into hardware performance than before, the provider can support its customers even more effectively. For example, AIOps helped the MSP resolve one of its customer's long-standing performance issues in a remote desktop environment. Detailed metrics and analytics revealed that it was related to an overly long CPU queue — something that could never have been identified before without considerable investigative and diagnostic work. The overall result is improved performance and availability for customers, together with improved more agile, efficient IT operations.

Breaking down silos and sharing global insights

One of the advantages of AIOps is its ability to bring together data from disparate sources and apply AI algorithms to enrich it with context—and minimize the volume of irrelevant alerts.

For example, a manufacturing firm in the UK depends on the availability and performance of its IT and networking infrastructure, which is located in data centers spanning the globe. Its locally managed monitoring tools were keeping individual systems up and running, but provided no single overview of the infrastructure as a whole.

For the fast-growing company, a centralized monitoring system became essential to bridging the gap between technical and service delivery teams by providing visibility across the worldwide infrastructure. An IT infrastructure monitoring and observability platform, with AIOps capabilities enabled the firm to consolidate views under a single pane of glass and customizable dashboards. Its analytics help the company avoid duplicating monitoring data, saving administrators time tracking down issues, and providing IT a single source of truth from a unified set of data.

With AIOps alert thresholds, support engineers on the front lines can identify and resolve issues before they reach a stage that requires escalation. In one instance, a support engineer received a service call about an air conditioning outage at a data center. Since AIOps features were used to track the temperature of all the company's servers, the early alert prevented the A/C outage from turning into a data center outage, which would have affected hundreds of people across the organization.

Gaining DevOps observability with machine learning anomaly detection

Like IT operations organizations, DevOps teams today are faced with managing rapid growth and complex infrastructures. Traditional static thresholds cannot offer the context and agility needed to manage these environments, so modern DevOps teams rely on advanced ML/AI algorithms.

Anomaly detection within an AIOps early warning system should provide context, meaningful alerts, illuminate patterns, and enable foresight and automation. It should be done automatically, without exposure to ML/AI algorithms and parameters.

AIOps can provide several benefits for anomaly detection for DevOps teams. It can predict issues before they occur, to prevent severe problems that can impact major initiatives. It also suppresses alerts on issues that don't require action, to minimize excess noise and alert fatigue. And it can help DevOps engineers troubleshoot engineers as they occur, helping them to determine if an issue is normal, and whether it was tied to a change in their environment.

Driving innovation and agility with DevOps

Like most business initiatives, DevOps is increasingly driven by cloud strategies. According to a recent study, nearly three quarters of global IT professionals believe that 95% of all public, private, and hybrid workloads will run in the cloud by 2025.[14]

All these forms of cloud technology enable you to deliver more value to customers faster. To better execute their cloud strategy, organizations have adopted a DevOps approach that embraces automation, monitoring and observability concepts.. AIOps should provide the full visibility of your IT infrastructure and workloads—whether in the data center or cloud—required to successfully manage a DevOps environment.

Automatically choosing the right anomaly detection algorithm

Several anomaly detection techniques are available, including Forests, Tensor-based, correlation-based, Neural Networks, Bayesian Networks, and deviations from association rules and frequent item sets. However, it's not always clear which algorithm is best to apply for a particular issue.

A modern monitoring and observability platform that incorporates AIOps capabilities should process data in a stream, keeping the system agile so it can quickly adjust and use the right algorithm. DevOps engineers should not need to become a data-scientist, but have the confidence that the AIOps implementation can do the hard work for them.

Extending the outlook beyond fixing issues

It's clear that AIOps enables IT professionals across the organization to gain insight and control that would never be possible using traditional monitoring and management approaches. But to compete effectively, you need to do more than simply respond to new problems—you need to be thinking several steps ahead of what you want your business to achieve, and how your technology will deliver those capabilities.

In our next chapter, we will discuss why AIOps is not only about issue mitigation, but also about continuous optimization. Using a failure prevention system, AIOps continually learns about your environment to optimize the technology stack, to position you to power digital transformation, innovate faster, and deliver the experiences that will keep your customers coming back.

[14] Evolution of IT Research Report [LogicMonitor], 2020

Chapter 9:

Future directions for AIOps: automated issue remediation

We've discussed how AIOps can help DevOps and IT operations teams achieve the business and technical outcomes they desire. AIOps capabilities are rapidly gaining traction and opening up new opportunities of insight across a variety of industries and environments. However, many of the most exciting benefits of AIOps are still just over the horizon. As AI algorithms and other technologies continue to mature, the ability of AIOps to augment human intelligence and accelerate action is setting up new possibilities.

In this chapter

In this chapter, we will look beyond current AIOps implementations to discuss its potential for driving continuous optimization. We will discuss:

- Introduce the basic requirements for a failure prevention system
- Integrating a failure prevention system into your existing tech stack
- The future state of AIOps: multiple data algorithms
- How remote work is driving AIOps adoption

Driving optimization with a failure prevention system

As we've introduced the capabilities of AIOps, most of our discussion has centered around its ability to help you discover potential problems earlier, and address issues before they impact your systems and processes. However, IT organizations are under constant pressure to continually optimize infrastructures that drive key business processes. The pace of change is accelerating and standing still is not an option if you're seeking to stay ahead of competitors and maintain a culture of innovation.

Using a failure prevention system, AIOps continually learns about your environment to optimize your technology stack, to put you in a strong position to achieve digital transformation, differentiate your organization by providing the best customer experiences, and maintain your competitive edge.

Defining and automating action

A failure prevention system builds on the data that is gathered and analyzed by AIOps, and increases the level of automated options, integration with other tools and systems, and proactive capabilities. For example, if you're on an IT team, it could give you the ability to:

- Define a specific action, such as execution of a custom script
- Set up a predefined action in response to an alert
- Automate those predefined actions in response to specific alert using a rules-based engine

For example, consider the use case we discussed earlier, where an SRE is alerted about database latency, and AIOps determines that the root cause was a change in configuration on a service that connects to the database. With a failure prevention system, the SRE could receive not only an alert on their phone that specifies the cause, but also get a dropdown menu with a list of predefined options—including the option to restore the service's previous configuration.

With a single click, the SRE could execute a script to revert the configuration and solve the database issue. An effective failure prevention system should support integration with an open-source provisioning and configuration management tool, such as Ancible, to executive even complex actions.

Nonstop learning about environments

As AIOps capabilities achieve broader adoption, AIOps solution vendors will acquire more data about their customers and apply these learnings to further enhance and automate capabilities. They could examine and analyze a vast array of big data, examine patterns, make predictions, and apply what they have learned to customers with similar environments.

For example, if an AIOps vendor discovers that a specific set of configuration changes can cause an Apache Web server to fail, they could deliver proactive alerts, or even execute scripts, to prevent customers from taking these actions.

Tracking the next wave of AIOps

The technologies that form the basis of AIOps are evolving rapidly, creating new possibilities for future AIOps solutions. There's no question that AIOps capabilities will only continue to expand in the months and years ahead.

AIOps is moving from one data type to multiple data type algorithms

One of the biggest trends in AIOps is applying its capabilities from one data type to multiple data types. The trend started where different probabilistic methods such as AI, machine learning, and statistical analysis were applied to a single data type that was either metrics, logs, or transactions.

In the near future, data scientists will be designing AI algorithms for multiple data sets together. It will look at the metric, log, and transaction data together, how they correlate, and what signals actually can be filtered out of all that noise to make troubleshooting issues faster. These algorithms for multiple data types will help you save time by enhancing early warning systems and filtering signals from noise more effectively.

Remote work is driving AI adoption

At the beginning of 2020, millions of workers around the world were forced to abandon their offices and work from home on short notice due to the coronavirus pandemic. For many companies, remote work will become the new normal. In fact, 74% of CFOs intend to shift some employees to remote work permanently.[15]

When it comes to AIOps, it doesn't matter where users are working. Once an algorithm is operationalized, its only job is to accept the input data, extract the intelligence, and output the optimized value. In remote environments, predicting performance that impacts a customer experience and employee productivity has become extremely important, because of the exponential increase in digital traffic.

As the workforce becomes more dispersed and remote, the data they generate is collected from more diverse locations—and has different properties. These disparate data streams are difficult to analyze manually due to the sheer volume of the data, which is where AI helps. AI can automate complex processing of disparate data sources and help you predict problems before they occur at an individual level, by detecting similar patterns in large volumes of data.

AIOps will become more embedded in observability platforms

Observability platforms look at metrics, dependencies, and logs, bringing them together to connect the dots between the different data types. Looking at this data gives you good observability across the customer experience, employee productivity, and digital infrastructure to understand how the business is performing. Incorporating AIOps and automation into these platforms reduces the time required to predict and fix problems before they impact your business. The usage of observability platforms in organizations is on the rise, along with the expectation for AIOps to become even more embedded in the near future.

Security and IT operations will be better integrated

To secure your business infrastructure and applications, the fundamental data is almost the same as IT operation data sets: the machine and user data flowing through your digital infrastructure. Security algorithms model the historical behavioral patterns and detect anomalies and deviations from those patterns in near real time. Using AI, this process could be further automated towards blocking bad actors in near real-time.

For example, suppose a hacker is trying to access or penetrate a firewall. Their attack could be detected by either a change in the volume of data, or a change in the location of the user that is trying to access it. Multiple features could be used to classify that particular access as either regular access, hacker access, or insecure access. Once that is detected, it could be handed over to the automation system to block the IP address of that particular region or that particular range.

The underlying data required to gather this intelligence is transactions, logs, and metrics,
but utilized by security teams. As IT operations and security teams work more closely together, they will improve their ability to not only detect problems in infrastructure performance but also prevent cybersecurity threats in near real time.

AIOps time to value will decrease

Many AIOps platforms have long setup times, consuming valuable DevOps resources for configuration. As AIOps vendors continue to enhance the capabilities of their platforms, they will increasingly embed better actionable insights and new proactive capabilities within their products—without significant setup costs. This advancement will set the foundation for future integrated self-healing systems.

Opening up the potential of business data

As the AI and ML algorithms that support AIOps gain new capabilities, they will enable you to take full advantage of the business intuition that is embedded in your organization's data—whether something's happening in the moment or patterns predicting the future.

Conclusion

LogicMonitor's SaaS-based platform seamlessly monitors everything from networks and infrastructure to applications and the cloud. A core part of that platform — AIOps enables teams to adopt a predictable, proactive approach to troubleshooting and monitoring. From data forecasting to dynamic thresholds, AIOps is designed to reduce MTTR and keep customers smiling. No matter how complex the environment, get meaningful context and correlation from LM Envision.

Automate more with AIOps

[15] CFO Actions in Response to COVID-19 [Gartner], Mar. 2020

Glossary

Artificial Intelligence

AI applies advanced analysis and logic-based techniques, including machine learning, to interpret events, support and automate decisions, and take actions

AIOps (Artificial Intelligence for IT Operations)

AIOps applies ML/AI algorithms to automatically detect anomalies in an infrastructure, application, or service, before they become problems, building on observability, automation, and an early warning system.

Analytics

Analytics supports a variety of different business intelligence (BI)- and application-related initiatives, analyzing information from a particular domain or applying BI to a specific content area.

Anomaly detection

Anomaly detection utilizes historical performance to generate an expected range for resources and dynamic thresholds, which is used to highlight and alert on anomalies that breach this range.

DevOps

DevOps focuses on rapid IT service delivery through the adoption of agile, lean practices in the context of a system-oriented approach.

Domain agnostic AIOps

Domain agnostic AIOps integrates with a variety of different services to collect data.

Domain centric AIOps

Domain centric AIOps primarily collects the data that is required to support AIOps on its own.

Dynamic thresholds

Dynamic thresholds are based on ML-based algorithms that automatically detect the normal performance range for any metric—whether it's a technical or business metric—and accurately alert based on values outside of this range that are considered anomalies.

Early warning system

An early warning system provides alerts about the most relevant issues in an environment, spotting issues and the warning signs that precede them, and triggering actions.

Failure prevention system

A failure prevention system builds on the data that is gathered and analyzed by AIOps, and increases the level of automated options, integration with other tools and systems, and proactive capabilities.

Forecasting

Data forecasting predicts future trends for monitored infrastructure, using past performance as the basis.

IT Operations

IT Operations are the people and management processes associated with IT service management to deliver the right set of services at the right quality and at competitive costs for customers.

Root cause analysis (RCA)

RCA is a systematic process for identifying the fundamental causes of problems or events and an approach for responding to them.

Virtual Machine (VM)

A VM is a software implementation of a hardware-like architecture, which executes predefined instructions in a fashion similar to a physical central processing unit (CPU).

651_AIOpsForMonitoring

About LogicMonitor®

LogicMonitor®'s SaaS-based observability and IT operations data collaboration platform helps ITOps, developers, MSPs and business leaders gain visibility into and predictability across the technologies that modern organizations depend on to deliver extraordinary employee and customer experiences. LogicMonitor seamlessly monitors everything from networks to applications to the cloud, empowering companies to focus less on troubleshooting and more on innovation. For more information, visit www.logicmonitor.com.

LogicMonitor