riverbed®

White Paper

What is AIOps?

Four Dimensions of Artificial Intelligence for IT Operations (AIOps)

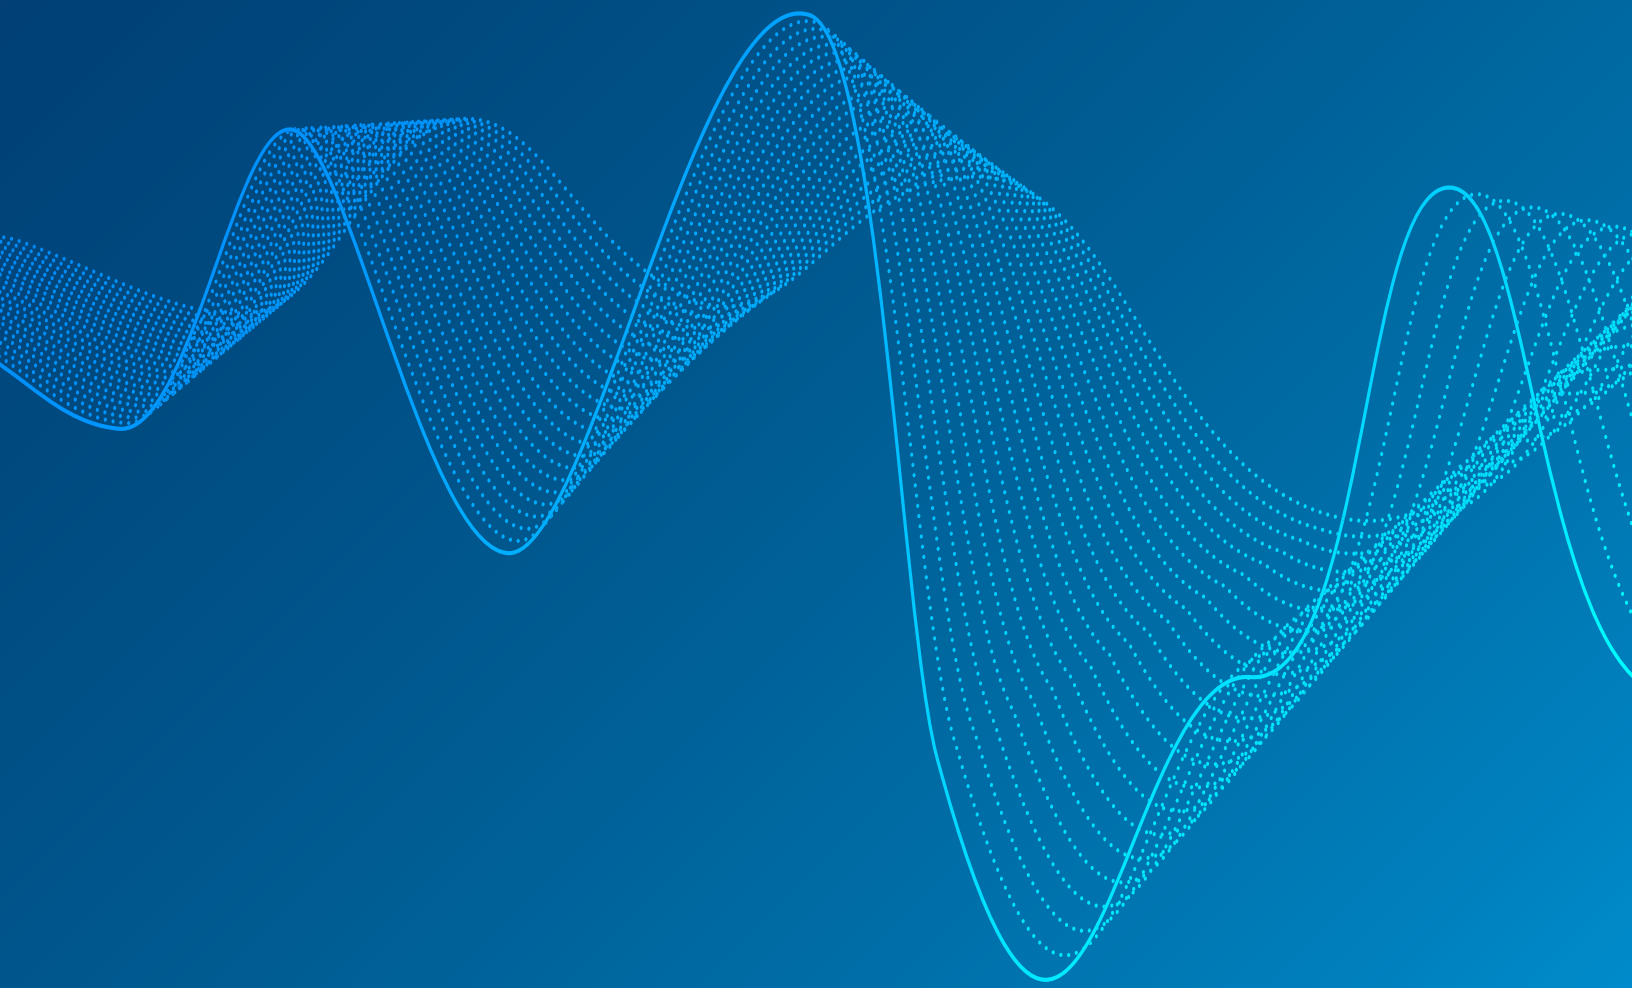

Contents

What is AIOps?

AIOps refers to the application of big data, machine learning, analytics and automation to IT Ops use cases in order to address this new need: to make sense of large quantities of mostly structured, specialized, cross-domain IT data.

"AIOps platforms combine big data and machine learning functionality to support all primary IT operations functions through the scalable ingestion and analysis of the ever-increasing volume, variety and velocity of data generated by IT. The platform enables the concurrent use of multiple data sources, data collection methods, and analytical and presentation technologies."

Gartner Market Guide for AIOps Platforms, Pankaj Prasad and Charley Rich, 12 November 2018

IT Ops teams can leverage machine learning and big data to drive continuous insights and automate remediation. This insight is used to drive incremental business value.

The capabilities of AIOps enable IT to:

- Reduce noise (such as false alarms) using clustering and pattern matching algorithms
- Determine causality, identifying the probable cause of incidents using machine learning (ML)
- Capture multi-variate anomalies to proactively detect abnormal conditions and behavior and relate them to business impact
- Detect trends that may result in outages before their impact is felt
- Drive the automation of low- to medium-risk recurring tasks
- Triage problems, helping prioritize them and offer actions that can be taken to resolve them

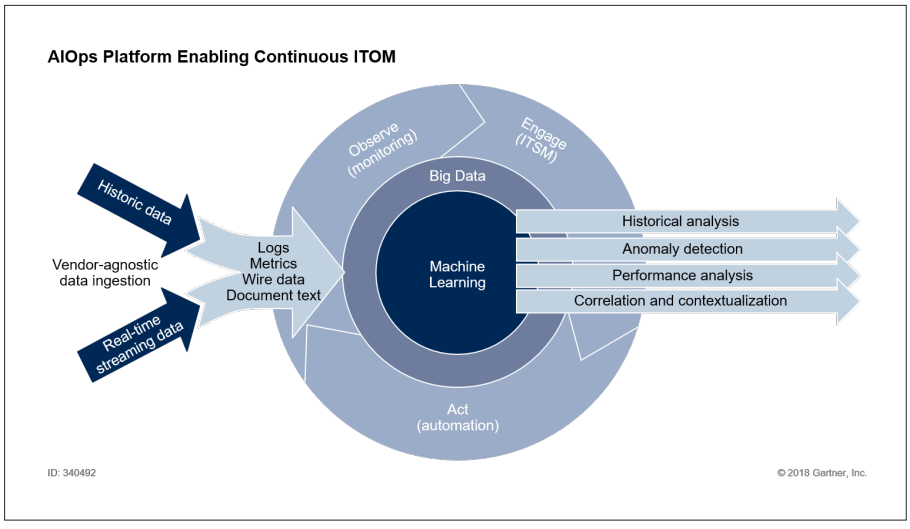

Figure 1: Ops Platform Enabling Continuous Insights Across IT Operations Management (ITOM)

Why AIOps?

In a word: Complexity

Complexity in every shape and form defines modern IT operations. From the complexity of multi-cloud and hybrid cloud to the mix of applications on prem, virtualized, and web to the complexity of mobile and globally distributed workforces, IT is faced with greater challenges to deliver the best user experience. Any performance issues in any application, its infrastructure, or underlying network can cost millions of dollars in lost revenue or employee productivity. Agility and cost efficiency are driving this complexity, yet at the same time, it exponentially increases the risks.

As complexity grows, so does the volume of data generated by IT that needs to be analyzed. Every platform generates its own set of metrics and user data. That's why machine learning is so critical – it can make sense of volumes of data too large for humans to timely analyze and gain actionable insights.

> "Rapid growth in data volumes (is) generated by the IT infrastructure and applications (two- to threefold increase per annum)."
>
> According to the Gartner Market Guide for AIOps Platforms, Pankaj Prasad and Charley Rich,12 November 2018

The key to unlocking the true power of AIOps is to leverage high-quality, structured data from multiple IT domains. Unfortunately, most IT tools operate in silos and rely on sampled transactions and low-resolution metrics—leaving large gaps for AIOps tools to extrapolate. This means that troubleshooters can find themselves chasing red herrings and automated remediation can be far too risky.

In this paper, we will explore four factors to consider when embarking on your AIOps journey:

1. Data
2. Context
3. Machine learning, visualization, and analytics
4. Automation

Dimension 1: You'll Need Data

Traditionally big data use cases have focused on the collection and analysis of unstructured data. Examples of unstructured data include email, video, audio, web pages, and social media messages, which is most of the data generated in the world today. Mining this data has its unique challenges because of the very nature of this human-generated data. However, IT has a unique leg up in that most IT data is generated by machines, and is therefore structured or at least formatted in a way that is well-defined and predictable. This creates immense possibility for AIOps initiatives to do more with IT big data than has been done before. **The key challenge is getting all the data—across multiple IT domains—into one place**. Data lakes have emerged as an effective way to do so and provide the foundation for AIOps.

The following are best practices for AIOps data:

- **Data granularity**: Low sampling rates of time-series data inevitably create distortions that can lead IT troubleshooters to the wrong conclusions or chasing after problems that don't exist. In signal processing, this distortion is known as aliasing effect. By collecting data at higher frequencies (i.e. 1-second intervals rather than 1-minute intervals), IT teams can more reliably identify surface anomalies when granular data is collected and stored for real-time and historic analysis.
- **Multiple data sources**: AIOps makes use of traditional monitoring sources across all IT functions. Examples include: transaction trace data, user and device data, application log data, observability data, infrastructure metrics, and network packet and flow data.
- **Data completeness**: At the end of the day, analytics are only as good as the data sets they are applied to. Incomplete or unreliable data sets will inevitably lead to wrong or inconclusive results. For instance, how would you identify a network latency issue impacting a cloud application that is limited to a specific ISP geography? Most monitoring tools would miss such an issue that only impacts a very limited geography but if it a senior executive is experiencing the issue, you will want to be proactively and immediately notified. It goes without saying: you won't find the needles in your haystacks if you're only examining a handful of hay from each stack.
- **Scalability**: The term "scalability" has evolved from the ability to support large enterprises to the ability to support increasingly large volumes of metrics and data. The reality of monitoring modern, cloud-based environments is a big data reality where the application environment changes rapidly and constantly.

18x more data collected, on average, from monitoring components and dependencies in container-based vs. more traditional monolith environments[1]

Dimension 2: Context is King

For information to be truly useful, it must be presented in context. Depending on the use case, context can be as granular as what network link is fully saturated to which users were affected by this issue and was there a significant business impact? Some analysts have gone so far as to say that more data can actually have a negative impact on the performance unless the data is delivered in a context that is actionable and relevant.

Take this example: If I am monitoring the response time of the web server in the cloud and it is slow, the normal response would be to double the capacity by spinning up a new instance using Kubernetes. Imagine that the first three times, slowdown is due to load and Kubernetes correctly spins up a new server. However, the fourth time the slowdown is due to another issue.

71% reported their IT performance data is not actionable[2]

If the AIOps system does not have sufficient context, it will continue to spin up web servers, each time doubling the capacity. This action, while correct the first three times, will not improve performance in the fourth scenario. It will, however, increase cloud costs. In this example, there is simply not deep enough visibility to understand the correct action to remediate the situation.

How can AIOps systems address this scenario? By collecting the metrics at a granular level in context, AIOps systems can better determine whether the automated remediation is appropriate or whether the anomaly should be surfaced for human review and intervention. In this case, the AIOps system may need visibility into the CPU utilization of the containers to identify whether additional CPU is needed. Or, the AIOps system might benefit from I/O metrics to show whether the I/O for the cloud infrastructure hit a bottleneck or limit that requires IT to upgrade their subscription to a more powerful option. Greater granularity in context can also assist IT to determine whether the issue is internal or external (such as an ISP problem) in order to engage the appropriate team as needed.

Note that if the AIOps system is self aware, it will recognize that the performance is not improving in this scenario and will cease spinning up new servers. Self-aware AIOps represents the future of AIOps.

Dimension 3: Surfacing Insights Through Visualization and ML

AIOps helps surface issues and causes beyond the usual suspects that you may not have thought to look for, point you in the right direction for faster troubleshooting and remediation, and proactively alert you when normal operating thresholds are breached.

Most of the work of AIOps is in surfacing patterns. Patterns can tell you quite a lot about your environment. They can point forward in time to potential issues that could develop into incidents or can be used to understand usage profiles that inform business or roadmap decisions. Looking backwards, patterns developing over time indicate potential root causes of issues and provide valuable clues for troubleshooting. When patterns show up in data, this usually points to something warranting further investigation.

To surface what's important, data scientists work with data in a number of ways:

- **Visualization**: Intuitive maps, graphs, and charts immediately surface key information and relationships among data. Using visualizations that capture every transaction rather than showing averages can better surface trends and highlight exceptions.

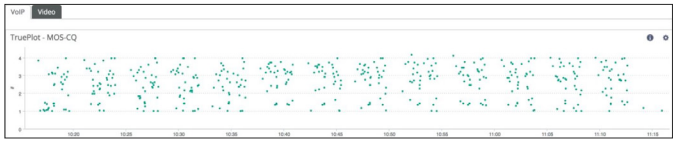

Figure 2: Visualization of performance clusters indicating a pattern to MOS scores.

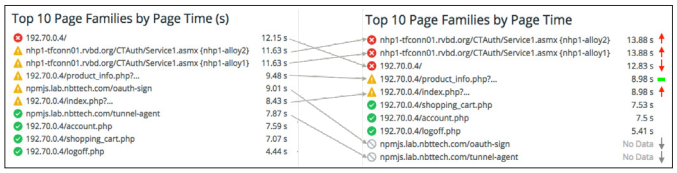

Figure 3: Ranked lists are an extremely valuable operational tool. In the report on the right, IT can identify the changes over time to the report as noted with arrows and greyed items that are no longer on the top 10 report.

- **Predictive analysis**: By applying analytics to data over time, AIOps is able to extract trends in the data and make predictions about future behavior. For example, predictive analysis reports can identify that in two days, you will exceed threshold for response time. The best practice is to include the confidence level around the prediction allowing IT teams to take action accordingly and prioritize issues.

- **Pattern recognition algorithms**: While some patterns are easily identified visually, others lie hidden inside vast amounts of data. The primary goal of machine learning is to automatically surface these patterns by looking for regularities in data sets. With unsupervised learning, the more granular and complete the training data set, the better the results you can expect.

- **Automated anomaly detection**: Anomaly detection is the ability to statistically learn what is normal and what is not for different sets of transactions. It proactively watches for transactions and metrics that are not behaving as they should and, if done right, takes seasonal variance into account.

- **Correlation**: By statistically correlating key performance indicators across all application tiers, the system can find groups of metrics on tiers that are spiking in tandem. This narrows down the set of possible culprits for a troubleshooter to investigate, which is especially helpful in chasing down performance issues in distributed containerized and microservices environments with thousands of nodes.

- **Event management**: This is one of the first areas that machine learning was applied to IT analytics. What's top of mind for IT today is to build more actionable context into alerts (events) and optimize the incident management workflow by incorporating automated remediation and cross-domain root cause analysis.

188% increase in processed metrics, events, and alerts over the last 12 months[3]

Dimension 4: Automation and Remediation

While the AIOps system can find meaningful correlations, correlation is not necessarily causation. Every infrastructure and every problem is different, and training neural networks to recognize general failure scenarios and root causes is a very difficult AI problem. There is no one magic set of tricks that just always works. This is where purely algorithmic approaches need to be bolstered with expert knowledge—embedded into the algorithms themselves—to improve outcomes for specific sets of problems.

In the infrastructure context, candidates for automation include re-routing network traffic to reduce congestion and free up bandwidth, spinning up additional cloud instances and concurrently expanding the SD-WAN fabric, and re-distribution of containerized workloads. For other sets of problems, the system may simply present prescriptive or actionable conclusions to the operator.

Automated remediation

Most IT organizations have runbook processes that specify the actions for IT to take in response to commonly expected problems. These processes clarify how to analyze the problem conditions, diagnose the issue, then fix it, based on a set of recovery actions. The recovery action is comprised of either manual steps or running scripts to clear the condition and verify proper operation.

These runbooks are ideal candidates for automated remediation but it's important to consider when a fix should be automatically applied.

Some key questions to ask are:

- How frequently does this problem occur?
- What is the impact/criticality when the problem occurs?
- Does the problem manifest itself in the same way each time or is there variability?
- How long does it take to fix? (For example, if it will cause a two-hour system reset, you may want to have a human examine the issue before proceeding.)

Once you have determined that the fix can and should be automated, AIOps systems can be enabled to support automated remediation. For audit purposes, the system can maintain a log of remediation actions taken, their success or failure, and who implemented them. Again, the AIOps system should be self aware and when automated remediation fails to correct the issue, it should escalate to the IT team.

This self-service capability results in a number of advantages for IT support teams:

- Improves user satisfaction
- Reduces trouble ticket volume and duration
- Improves first level resolution rates

The Final Link: Human Intelligence

It's realistically impossible to have a truly complete dataset–a full understanding of a complex, ever-changing infrastructure that is inevitably affected by factors outside of what can be monitored. That missing piece often comes from the human operator. Remember the old story about the website that was mysteriously crashing after hours? It turned out that the cleaning staff was regularly unplugging the servers to plug in the vacuum! Some missing piece of information in the operator's head could instantly make that causal link between those website crashes and the cleaning schedule even if the AIOps product cannot.

That's one reason why you won't always want to let the system decide to execute remedial actions on its own. While AIOps can have the built-in intelligence to suggest specific remedial actions that are often appropriate—which is great for taking a lot of cognitive burden off of operators—having a human in the loop to make that final call is just a wise thing to do.

AIOps in Practice

The power of AIOps is in analyzing IT big data and taking action faster than humanly possible–to drive better business outcomes. Today's IT organizations can benefit from AIOps, leveraging machine learning and visualizations across extremely large, cross-domain datasets. Using AIOps, IT can accelerate root cause analysis, automate remediation, and ultimately drive better business outcomes. This is not hype; it is happening today across industries on a large scale.

Forward-thinking IT organizations are applying automation and remediation to create a self-learning and self-healing system that addresses issues before anyone is ever aware they exist. Human intelligence can and will continue to play an important role even as automation is rolled out: there are certain scenarios you simply cannot automatically resolve and others can benefit from embedding human intelligence into the algorithms.

The AIOps journey is only beginning.

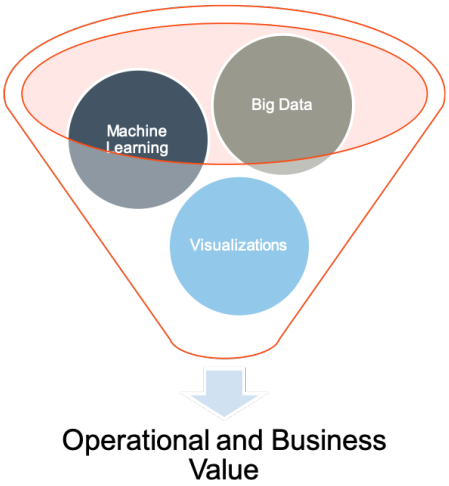

Sources:

[1] Digital Enterprise Journal, 17 Areas Shaping the IT Operations Market in 2018.

[2] ibid.

[3] ibid.