White paper

virtana

A Guide to AIOps

The Future of IT Ops

virtana

Over the last few years many IT monitoring solutions haven't kept pace with the advances of modern IT infrastructures. Many organizations still rely on disparate and siloed monitoring tools built on legacy framework that present a fragmented view of IT operations. It is increasingly clear that tools developed to keep a single system or a small cluster running are no longer sufficient in today's highly distributed, complex environment.

If you are in IT Ops or DevOps, hardly a day goes by without someone mentioning AIOps. There are a few who think AIOps can replace IT Ops tools today. Others debate this, saying that AIOps is still a nascent field, and it will take a few more years until we see a full-fledged AIOps platform for IT operations management. But there's always been a lot of confusion on how AIOps really works.

AIOps came about because a new set of tools that offer service assurance rather than merely offering high availability are needed. These tools will be able to monitor and manage all the components of the IT infrastructure, not just a few selected components to achieve the goal of preventing service interruptions that have a negative impact on users. AIOps helps you to clearly understand how performance and availability degradations impact overall service delivery while uncovering trends and patterns that help improve troubleshooting, service reliability and quality across the technology stack.

What is AIOps?

Defined by Gartner, artificial intelligence for IT operations (AIOps) platforms utilize big data, modern machine learning (ML) and other advanced analytics technologies to directly and indirectly enhance IT operations (monitoring, automation and service desk) functions with proactive, personal and dynamic insight. When Gartner coined the term 'AIOps' a few years ago, it stood for "algorithmic IT operations." But because AI was in the acronym, confusion ensued, and, eventually, Gartner changed their usage of the phrase to be "artificial intelligence for IT operations." First-generation AIOps tools performed real-time analysis on mass quantities of event data and inferred probable root causes based on data analyzed from previous issues. Traditional AIOps tools are not service-centric and have no concept of topology. They rely solely on processed event data and suffer from blind spots that come from having no visibility into other data types, e.g., metrics, dependency data, streaming data, logs and other types of machine data.

Second-generation AIOps tools are beginning to emerge, and the key difference is that these solutions collectmore than just events. The data they collect includessome combination of events, metrics, logs, streamingdata, dependency data and more, eliminating the key problem experienced by AIOps tools, i.e., limited visibility and context due to the lack of cardinality in the data they analyze. It enables vendors to inform ML algorithms with explicit topology, which makes a vast difference in detecting and isolating issues with certainty. This provides unprecedented context and unprecedented acceleration of problem resolution. Today AIOps addresses key areas, including data collection and storage, analytical engines (real time and deep), visualization/UI, and integration with other applications. AIOps is seen as a way to reduce MTTR and to sort through noise with algorithms, visualization, retrospective analytics and dashboards.

Deconstruct Different Layers of AIOps

At a high level, AIOps tools do two things: they collect data and they analyze data. They do this in the interest of accelerating problem resolution in IT operations. But reducing AIOps to just these two broad areas completely ignores many of the foundational elements of true AIOps solutions — especially ones that really seek to successfully leverage AI in order to streamline IT operations. We've identified nine distinct capabilities — with a bottom-up approach — starting with data collection, then moving up through the layers to self-learning capabilities.

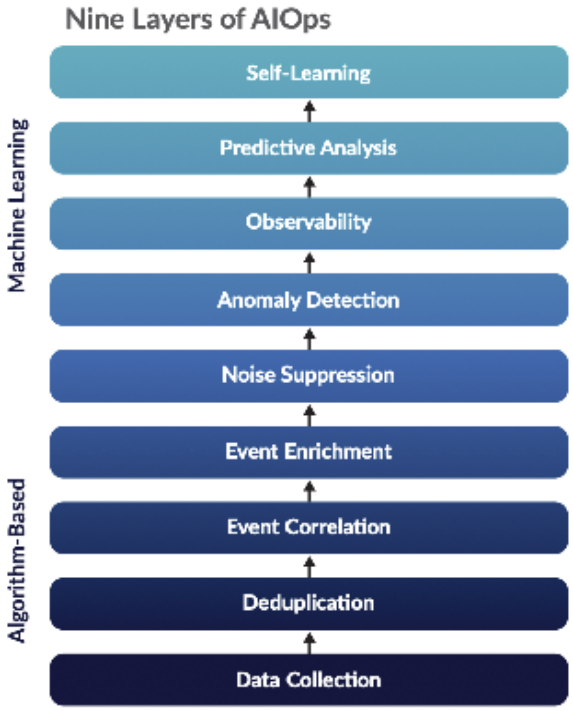

Data Collection - This is the key foundational component that fuels complex machine learning algorithms to analyze and see patterns that IT Ops teams, in general, wouldn't find. If you are a large organization managing thousands of devices and millions of nodes globally, it is highly unlikely that you could manually sift through all available complex machine data in your environment, including metrics, events, streaming data, traces and log data, and predict upcoming application or system failure within any given time. A flexible, comprehensive data collection mechanism is what any AIOps tool needs. But IT Ops teams are realizing that theirexisting tools have limited capabilities and that they need an expansive breadth of data to create meaningful machine learning models.

Contrary to what many stand-alone AIOps vendors claim, adding intelligence on top of logs or events alone is not sufficient to create trustworthy automation. For robust self-healing, these tools must be able to stream all data types from different sources to provide enough context to be trustworthy. For instance, Virtana provides a deep data collection mechanism that enables IT Ops to collect streaming data in real time and gain insights on how to optimize and improve IT performance.

Deduplication - Automatic deduplication is the next critical step when you are processing millions of oncoming events from different data sources at a large scale. For many admins managing hybrid IT environments with some combination of on-premises and public/private cloud infrastructure, reducing event noise automatically saves a lot of time and improves mean time to resolution (MTTR).

Event Correlation and Enrichment - Most monitoring solutions handle event storms by reducing and correlating events across different data sources using statistical analysis. Some stand-alone AIOps solutions approach event correlation with pattern-recognition machine learning (ML) algorithms, but without training the ML algorithms using large sets of contextual data, they fall short of adding much value to users who tend to look for anomalies or insights from the IT environment. For the person or team responsible for the event source, the impact on the business and the initial steps for triage are critical bits of information generally lacking from alerts. You can leverage model-based data and context- aware relationships from your IT infrastructure and use that to suppress and enrich events through AIOps.

Noise Suppression - IT Ops teams always look for actionable alerts that enable them to become more efficient. But one common hurdle is excessive alert noise and fatigue, which can be a hurdle for IT admins who usually seek faster resolution time. Some AIOps solutions deliver intelligent event suppression, enabling admins to selectively discard

and reduce event noise — making it easier for them to scale device monitoring and zero in on important problems. It works if you are in large application environments where you have a lot of noisy event sources.

Anomaly Detection - This layer of AIOps is where the technical complexity goes up a notch. For instance, if you are deploying code across your environment, you would want to know the potential service impact. By having advanced machine learning algorithms, you can save development costs by automating the manual process of sifting through the logs and metrics to find out if there are any anomalies or regressions.

Observability - One of the key challenges for stand-alone AIOps solutions is the lack of a tremendous amount of raw machine data to train their machine learning algorithms for better observability. For instance, the analysis of one or two types of data (such as metrics, events, logs, tracing, etc.) won't provide you with necessary insights to debug a hybrid cloud application. You need machine data from all sources/ types to fully realize the service context, and machine learning can help you make sense of the mix of high- cardinality data. When you start with a robust and contextual data collection mechanism from different sources as your foundation stone, it enables observability and helps you to act on critical insights that allow you to analyze root cause and speed up troubleshooting and the resolution process.

Predictive Analysis - Given the dynamic nature of IT, where relationships between services and IT infrastructure are constantly changing, proactively preventing impacts to critical business services is never easy. The most useful machine learning insights are informed by real-time infrastructure modeling data as well as large datasets from all other sources. You can stay ahead of the curve by uncovering model-informed trends and patterns that help optimize your IT resources and plan capacity more effectively.

Self-Learning - Many AIOps vendors are on a path to accelerate machine learning effectiveness with real-time, dynamic models of end-to-end IT services and applications. With supervised or unsupervised training of ML algorithms, including deep learning, AIOps solutions can effectively forecast system or performance failures even before they happen. One potential upside of expanding AI capabilities beyond these nine layers as AI/ML algorithms become more advanced is that it might even push the limits of AIOps tools to self-heal whenever there is an issue in IT infrastructure. But, for now, we are just in the early stages of adopting and using AIOps capabilities for IT operations management.

Can AIOps Replace Monitoring?

No. Stand-alone (first-generation) AIOps tools were created to address a specific problem: too many monitoring tools. Medium and large enterprises typically have in excess of 30 monitoring tools. The idea of AIOps was to overlay a technology that could ingest events from all the monitoring tools, correlate them, and give you inferred insights. But, one of the biggest challenges is that this approach relies completely on pattern matching. For this type of tool to precisely identify the root cause of an IT issue, it must have seen that exact issue with the exact same fingerprint (the "pattern") some undetermined number of times. In today's complex, dynamic environments, these issues rarely have the same fingerprint. So, what early adopters of stand-alone AIOps tools have learned is that they must endure countless disruptions/outages before the tools begin to provide real value.

With the help of second-generation AIOps tools, unified monitoring tools can inform ML algorithms with topology. In other words, it's telling the algorithms exactly how various systems are connected and dependent upon each other to deliver an application or IT service — not leaving it for the algorithms to infer. This dramatically changes a tool's capacity for precisely pinpointing the root cause of an IT issue.

Collect More Data for Better Context

Some AIOps tools helps in aggregating and contextualizing data to provide timely insights for IT Ops teams. Over time, you can validate the results of AI/ML algorithms and test the reliability of these forecasts. AIOps tools can help in analyzing unstructured data in order to identify higher-level correlations that traditional IT monitoring tools wouldn't be capable of. But IT Ops teams often struggle with drawing meaningful insights from this data due to lack of context. They invariably end up doing data manipulation or aggregation by setting up models of data ingestion without thinking about the business first.

The success of AIOps and other automation methods relies on the quality of the data. Good, representative data is what any machine needs to operate accurately and perform the task that meets the needs of the business. Contextualized data, when taken together from a variety of sources, provides an accurate view of what's really going on in the business.

Most IT infrastructure monitoring platforms have relied on two key sources of data: performance metric data and infrastructure metadata. Intelligent application and service monitoring solutions like Virtana integrate with Splunk, Moogsoft, BigPanda, etc. to simplify this wave of data by sorting through the noise to highlight the key events by enabling cross-functional collaboration. This also helps in finding higher-level, seasonal trends beyond usual monitoring metrics and indicators.

An example of an application built on an AIOps platform that spans multiple ITOM functions is an actionable, comprehensive feedback loop for a DevOps-delivered application to drive its continuous improvement. Some enterprise DevOps teams have done exactly this, building applications of this scope for a given application that include data from monitoring, automation, service desk and application development tools using AIOps platform. The key to the decision to use an AIOps platform is that AIOps platforms uniquely provide more than just a method for gaining visibility into all the activities associated with an application's creation, performance and evolution. Importantly, they also add the capability for both machines and people to learn from the behavior of the people and systems involved.

Analyze Complex Machine Data

AIOps tools can help analyze unstructured data to identify higher-level correlations that traditional IT monitoring tools wouldn't be capable of. IT Ops teams are constantly bombarded with complex machine data, and to better monitor and troubleshoot IT environments, you must analyze and gain an accurate understanding of this data to evaluate how systems are performing and proactively resolve any issues that may occur. However, traditional monitoring tools don't help much in navigating through this data while managing the added intricacies of serverless architecture, microservices, containers, and other technologies. AIOps tools perform better with more data, so they should seamlessly connect to different data sources in your IT environment and ingest the data that they generate. Typically, the data sources would include the monitoring data, metrics data, application logs, model data, etc. These data sources are continuously generating huge volumes of data, in both structured and unstructured form, and hidden within this data are insights that would help in proactive IT infrastructure management. So, one of the key elements of AIOps should be a big data platform for managing the data that is being ingested.

AI and machine learning algorithms access both the historical data as well as real-time streaming data. The real differentiation that AIOps tools bring is the ability to continuously learn and optimize function based on new data. For instance, if a monitoring tool alerts that the increased CPU usage is due to an increased number of connections, Kubernetes can spin up the additional app instances and use the load balancing to distribute the users and reduce the load. AIOps capabilities in advanced IT infrastructure monitoring solutions automates routine DevOps tasks, enabling the machine learning model to launch it under certain conditions and deal with the issues preemptively, before downtime occurs.

AIOps tools work to reduce alert noise when event storms occur by performing real-time analysis on mass quantities of event data and inferring probable root causes based on data analyzed from previous issues. But, AIOps tools are not service-centric and have no concept of real-time models. They are not designed to tell the user how issues with a given system may affect IT services or applications. Virtana combines the probabilistic AIOps event correlations with root-cause analysis informed by definitive service- dependency models, eliminating guesswork when investigating IT incidents.

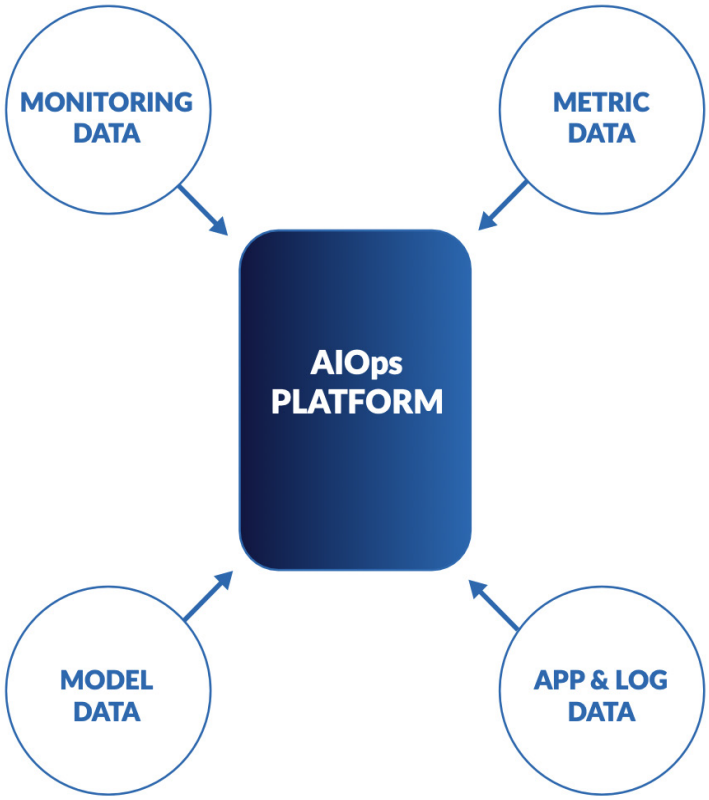

What is the Future of AIOps?

AIOps platform technologies have been most frequently adopted in support of availability and performance monitoring efforts. This is due to a number of factors, most notably the need of monitoring teams to rapidly perform often highly complex diagnostic tasks that AIOps technologies are ideally suited for. However, as IT operations tasks become increasingly automated, and roles and responsibilities continue to converge the work of analysis becomes a growing portion of all IT operations functions. This convergence in turn results in a growing need for AIOps platform capabilities that both AIOps-platform-focused and domain- centric vendors will continue to work to fulfill. Domain-centric vendors will continue to add AIOps platform technologies in various forms in a bid to become the dominant platform vendor, and current AIOps- platform-focused vendors will continue to add capabilities that make them an increasingly viable alternative to domain-centric tooling.

Most industry analysts agree that AIOps will be a key element of IT technology stacks for the foreseeable future. Many industry analysts also agree that the pioneers of AIOps (the first-generation, or stand-alone, AIOps tools) will have increasingly diminished value as event collection must be augmented with higher-cardinality data to make the solutions viable in complex, modern IT environments. Second-generation AIOps vendors like Virtana are delivering a new level of intelligent analytics capabilities for all data types, including metrics, dependency data, events and streaming data, providing unprecedented context and unprecedented acceleration of problem resolution.

How Virtana Helps

Virtana gives full visibility into IT service relationships and dependencies through real-time modeling that stand-alone AIOps solutions can't provide. Virtana can augment insufficient log and event data used for AIOps correlation with rich metrics from every system constituting every IT service. As a leader in intelligent application and service monitoring, Virtana can provide machine learning insights informed by real-time model data, as well as all other data types. This creates an unprecedented capability to visualize incidents, forecast trends and detect issues before the business is impacted.

Virtana has introduced full-stack monitoring with AIOps. This means eliminating the number one problem AIOps tools have experienced thus far — limited visibility and context due to the lack of cardinality in the data they're analyzing. Virtana is delivering a new level of AIOps analytics capabilities for all data types, including metrics, dependency data, events and streaming data. This provides unprecedented context and unprecedented acceleration of problem resolution.

For more information on how you can utilize Virtana Service Observability to unify your observability practices across legacy and modern IT environments, contact Virtana to set up a demo. To learn more about AIOps, download the Gartner Market Guide for AIOps Platforms here.

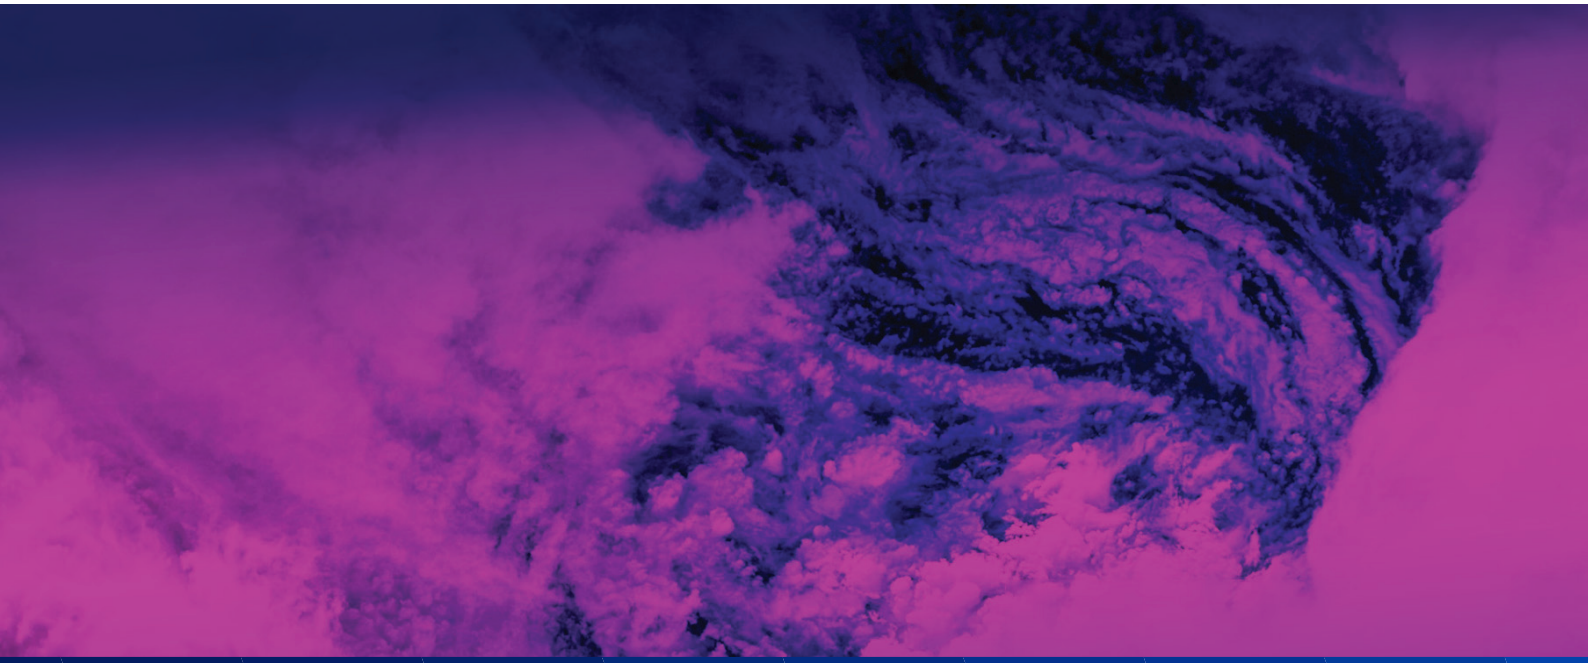

info@virtana.com | +1-408-579-4000 | virtana.com