Market Guide for AIOps Platforms

Published 12 November 2018 - ID G00340492 - 18 min read

By Analysts Pankaj Prasad, Charley Rich

AIOps platforms enhance IT operations through greater insights by combining big data, machine learning and visualization. I&O leaders should initiate AIOps deployment to refine performance analysis today and augment to IT service management and automation over the next two to five years.

Overview

Key Findings

- AIOps is getting entrenched in enterprises predominantly for IT operations, while some of the more mature organizations are using the technology to provide insights to business leaders.
- AIOps skills and IT operations maturity are the usual inhibitors in ensuring quick time to value when using these tools, followed by data quality as an emerging challenge for some of the more mature deployments.
- Enterprises adopting AIOps platforms use it to enhance and, occasionally, augment classical application performance monitoring (APM) and network performance monitoring and diagnostics (NPMD) tools.
- Vendors are developing strategies to use machine learning — the primary technology within AIOps — to analyze data challenges for IT operations across the three dimensions of volume, variety and velocity. At the same time, they are building specialization across both data storage and AI practices.

Recommendations

I&O leaders responsible for optimizing IT operations should:

- Deploy AIOps by adopting an incremental approach that starts with historical data, and progress to the use of streaming data, aligned with a continuously improving IT operations maturity.
- Select platforms that enable comprehensive insight into past and present states of IT systems by identifying AIOps platforms that are capable of ingesting and providing access to text and metric data.

- Deepen their IT operations team's analytical skills by selecting tools that support the ability to incrementally deploy the four phases of IT-operations-oriented machine learning: descriptive, diagnostic, proactive capabilities and root cause analysis to help avoid high-severity outages.

Market Definition

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.

Market Description

AIOps can enhance a broad range of IT operations processes and tasks, including performance analysis, anomaly detection, event correlation and analysis, IT service management and automation.

Their central function is:

- Ingesting data from multiple sources agnostic to source or vendor
- Enabling data analytics at two points:
  - Real-time analysis at the point of ingestion
  - Historical analysis of stored data
- Providing access to the data
- Storing the acquired data
- Using machine learning
- Initiating an action or next step based on the result of analysis

The goal of the analytics effort is the discovery of patterns — novel elements used to look forward in time to predict possible incidents and emerging usage profiles — and to look backward in time to determine the root causes of current system behaviors (see Figure 1).

Figure 1. AIOps Platform Enabling Continuous Insights Across IT Operations Management (ITOM)

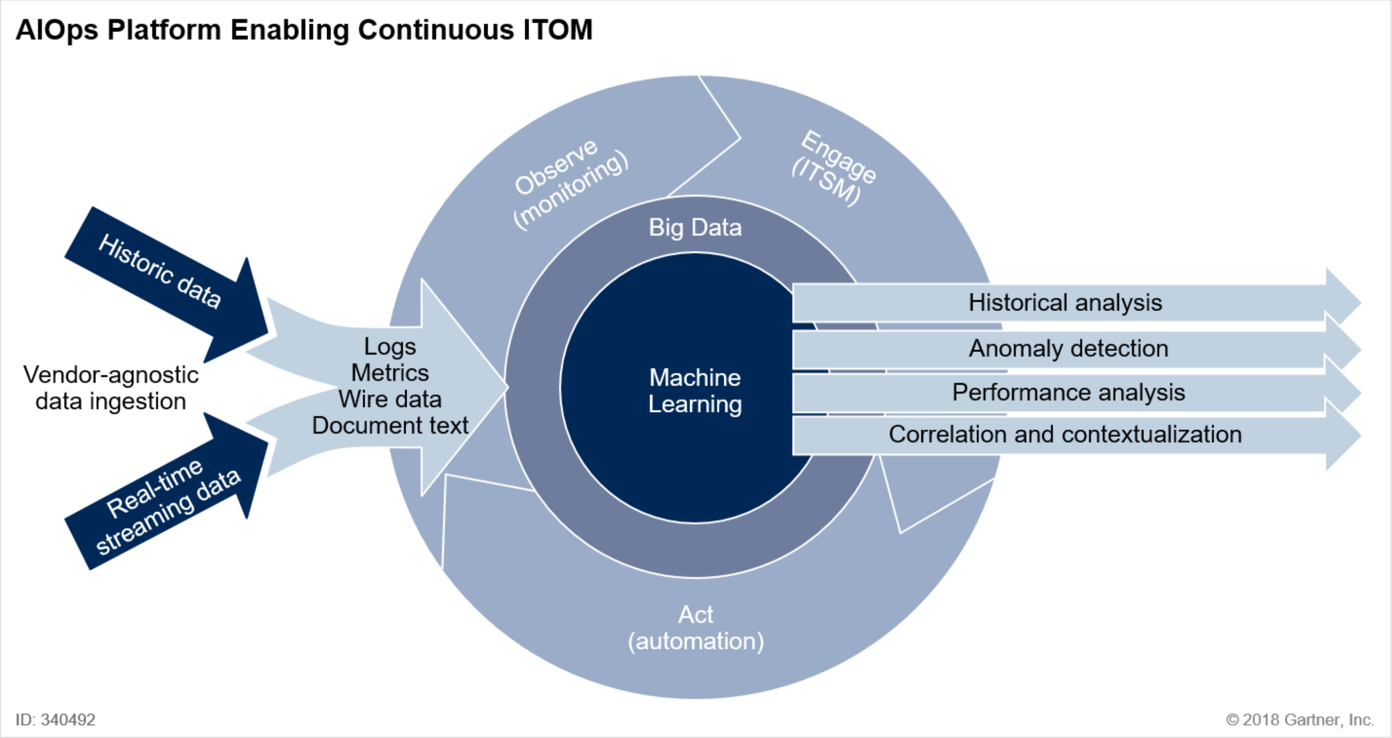

Source: Gartner (November 2018)

Market Direction

AI technology has influenced the evolution of ITOM intermittently over the past two decades, and AIOps platforms are only the most recent example of that influence. IT operations is challenged by the opposing forces of cost reduction on one hand and increasing operations complexity on the other. The complexity can be defined across the three dimensions of volume, variety and velocity as:

- Rapid growth in data volumes generated by the IT infrastructure and applications (two- to three-fold increase per annum)
- The increasing variety of data types generated by machines and humans (for example, metrics, logs, wire data and documents [knowledge management])
- The increasing velocity at which data is generated as well as the increasing rate of change within IT architectures due to the adoption of cloud-native or other ephemeral architectures

A trade-off in any of these dimensions will prove costly given the insights required by a modern business. Existing monitoring tools are stressed when dealing with high volume, variety and velocity of data. More importantly, monitoring tools do not cut across the multiple data types required for extracting useful insights. For example, the business needs enormous amounts of data that cuts across infrastructure and application metrics, customer sentiment data, business transaction data, sensor telemetry, and logs from various systems for additional insights.

Non-IT groups like line of business owners and teams that sit outside IT operations (such as application developers and DevOps) are increasingly showing interest in AIOps technologies to surface insights across a multitude of datasets (see “Artificial Intelligence for IT Operations

Delivers Improved Business Outcomes”). In some cases, security and IT operations teams are exploring opportunities to leverage a common platform (see “Align NetOps and SecOps Tool Objectives With Shared Use Cases”). The performance and maturity of the AIOps platform toward enabling the multiple use cases across IT and security operations have been primary inhibitors against a common platform deployment.

Further, the speed with which IT needs to act is also increasing due to digital business, hence the need for tools that can help:

- Reduce noise (for example, in the form of false alarms or redundant events)
- Provide better causality, which helps identify probable cause of incidents
- Capture anomalies that go beyond static thresholds to proactively detect abnormal conditions
- Extrapolate future events to prevent potential breakdowns
- Initiate action to resolve a problem (either directly or via integration)

To date, AIOps functionality has primarily been used in support of IT operations processes that enable the monitoring or observation of IT infrastructure, application behavior or digital experience. AIOps platform investments have almost always been justified on the basis of their ability to decrease mean time to problem resolution. And they have been justified regardless of whether this takes the form of using machine learning to deduplicate events in an event management context or to analyze application log data in conjunction with bytecode-instrumentation-based or distributed tracing data in an APM context.

AIOps platforms are expanding the range of data types they are capable of ingesting. In particular, vendors that supported only the ingestion of log data in the past are now expanding their scope to include metric and wire data.

Therefore, given both supply- and demand-side trends and technical differences, Gartner anticipates that, over the next five years, wide-scope AIOps platforms will become the de facto form-factor for the delivery of AIOps functionality as opposed to AIOps functionality embedded in a monitoring tool like APM, NPMD or ITIM (see “Deliver Cross-Domain Analysis and Visibility With AIOps and Digital Experience Monitoring”).

Gartner clients have demonstrated increasing interest in using AIOps functionality to improve engagement with incidents and problems by applying big data and machine learning to trouble ticketing to analyze the effectiveness of the service desk. [1], [2]

IT organizations have also started exploring AIOps in a DevOps context as part of the continuous integration/continuous delivery (CI/CD) cycle to predict potential problems prior to deployment and to detect potential security issues [3] (see “Market Guide for Continuous Configuration Automation Tools”).

AIOps analysis is expanding beyond its initial usage as a better solution for event correlation and analysis in IT operations. I&O leaders are beginning to focus on use cases beyond the realm of IT operations. As an example, since January 2018, Gartner clients have expressed interest in designing dashboards showing real-time analysis of customer satisfaction, the order process and business health. [4] The goal in this case is to present line of business owners with real-time insights into the impact of IT on business, keeping them informed and enabling them to make decisions based on relevant data.

Gartner believes that AIOps will evolve into a bidirectional solution that not only ingests data for analysis, but also initiates actions based on its analysis. These actions, most likely via integration to other ITOM and ITSM tools, will take several forms, including:

- Alerting
- Problem triage
- CMDB population
- Run book automation
- Application release orchestration

AIOps tools show a “right-shift” across the four stages of monitoring — data acquisition, aggregation, analysis and action (see Figure 2) — with their core capabilities at data aggregation and analysis. As the technology matures further, users will be able to leverage proactive advice from the platform, enabling the action stage.

With increasing instrumentation in modern applications, data acquisition as a native application capability is leveraged by some organizations. In addition, some users leverage open-source technologies for the data acquisition function, thereby bypassing APM as a specialized domain monitoring tool and using AIOps as the primary source for the monitoring function.

The debate regarding monitoring tools versus AIOps has just begun, and it will likely grow. Nevertheless, in the long run, monitoring tools will exist for the domain specialist, whereas, for an IT operations generalist, the primary go-to tool will be AIOps.

Figure 2. Four Stages of Monitoring

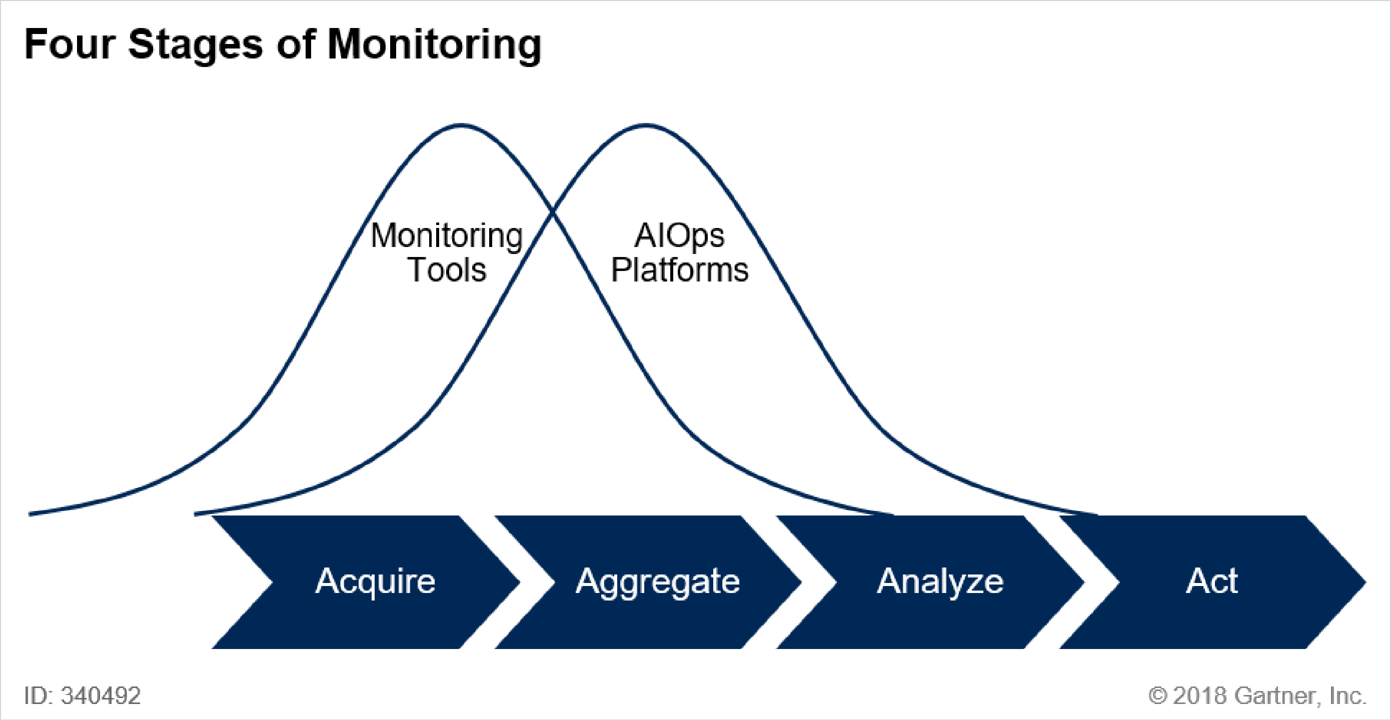

Source: Gartner (November 2018)

Market Analysis

To date, few vendors offer comprehensive, integrated AIOps platforms. Many vendors do, however, offer a wide range of AIOps capabilities, subsets of which are integrated with one another. To get a clearer picture of how the market is evolving and where vendors are positioned with regard to one another, Gartner has divided currently available AIOps capabilities into two major categories across data management and analytical outcomes:

Data Ingestion and Handling

- **Historical and streaming data management —** Software or appliances that allow for the ingestion, indexing and persisted storage of log data, wire data, metrics and document data (see Note 2). The resulting databases are mostly unstructured or polystructured, while the stored datasets accumulate in high volumes, change with high velocity and are implicitly structured according to highly varied formats. This historical data management functionality can be called "big data management." To provide value under the IT operations use case, the tool must also present data in time scales perceived by a human user as real time, delivering data directly at the point of ingestion without requiring access to a persisted database. It must provide a coherent analysis across multiple streams of real-time and historical data.

Analytical Outcomes

- **Basic and advanced statistical analysis—** A combination of univariate and multivariate analysis, including the use of correlation, clustering, classifying and extrapolation on metrics captured across IT entities as well as for curating data at source.
- **Automated pattern discovery and prediction —** Use of historical or streaming data of one or more of the types mentioned above, to elicit mathematical or structural patterns that describe

novel correlations that may be inferred from, but are not immediately present in, the datasets themselves. These patterns may then be used to go forward in time and predict incidents with varying degrees of probability.

- **Anomaly detection —** Using the patterns discovered by the previous components to first determine what constitutes normal system behavior, and then to discern departures from that normal system behavior.
- **Root cause determination —** Pruning down the network of correlations established by the automated pattern discovery and prediction component to isolate those links of dependency that represent genuine causal relationships in the sense of providing recipes for effective intervention.
- **Prescriptive advice —** Performing triage on problems, classifying them into known categories. It may then mine stores of prior solutions, analyzing these for applicability and offering them in a prioritized form for usage of remediation. Eventually, these will use a closed-loop approach and enable voting on their effectiveness after they are utilized.
- **Topology—** For the patterns AIOps detects to be relevant and actionable, a context must be placed around the data ingested. That context is topology. Without the context and de facto constraint of topology, the patterns detected, while valid, may be unhelpful and distracting. Deriving patterns from data within a topology will reduce the number of patterns, establish relevancy and illustrate hidden dependencies. Using topology as part of causality determination can greatly increase its accuracy and effectiveness. Capturing where events occurred and their up and downstream dependencies using graph and bottleneck analysis can provide great insight on where to focus remediation efforts.

There is some confusion in the market concerning whether AIOps will replace domain-centric monitoring tools such as APM, NPMD, ITIM and DEM (see “Hype Cycle for IT Performance Analysis, 2018”). AIOps will not replace monitoring tools, rather it will provide enhanced analytics and more actionable data. Domain-centric monitoring tools will continue to exist providing data capture, analysis and visualization of their domains for the specialist. However, they will forward their data streams to an AIOps platform, acting as a lens where the data will be focused into a single, coherent cross-domain analysis (see “Deliver Cross-Domain Analysis and Visibility With AIOps and Digital Experience Monitoring”).

As the market evolves, Gartner has observed evolving AIOps capabilities across various dimensions:

- Vendors going to market with a data-source-agnostic AIOps platform. These products tend to be generic and cater to the broadest use cases.
- Vendors that have the key components, but tend to have a restricted set of data sources. These vendors are typically focused on one domain (for example, network, endpoint systems and

APM), or are selective about data types like alert streams from other tools. Such tools tend to have a restricted set of use cases, targeted at a certain segment of IT operations.

- Some vendors with existing monitoring solutions limit data sources to their own monitoring products or extend to a limited partner ecosystem. This is again a case where the target audience is limited to those with the right mix of data sources.
- Some open-source projects enable users to assemble their own AIOps platforms by offering tools for data ingest, a big data platform, ML and a visualization layer. End users can mix and match the components from multiple providers.

AIOps platforms add important capabilities beyond what a monitoring tool with embedded AIOps can provide. The AIOps platforms free from the implied lock-in to a static data model expressed in monitoring tools is able to capture the patterns, anomalies and causal structures in the data itself. Monitoring tools may miss these features in the data as they force it into their predetermined models.

Representative Vendors

The vendors listed in this Market Guide do not imply an exhaustive list. This section is intended to provide more understanding of the market and its offerings.

Market Introduction

AIOps platform vendors have a broad range of capabilities that continues to grow. Vendors differ in their data-ingest and out-of-the-box use cases made available with minimal configuration.

In Table 1, we have provided a representative, sample list of vendors providing AIOps platform functionality.

Table 1: Representative Vendors

Vendors ↓	**Year Founded** ↓	**Headquarters** ↓	**Website** ↓
Anodot	2014	Israel	https://www.anodot.com/
BigPanda	2014	United States	https://www.bigpanda.io
BMC	1980	United States	https://www.bmc.com/
Brains Technology	2008	Japan	https://www.brains-tech.co.jp/en/
CA Technologies	1974	United States	https://www.ca.com/us.html

Vendors ↓	Year Founded ↓	Headquarters ↓	Website ↓
Devo (Logtrust)	2011	United States	https://www.devo.com/
Elastic	2012	United States	https://www.elastic.co/
Evolven	2007	United States	https://www.evolven.com/
FixStream	2013	United States	https://fixstream.com/
IBM	1911	United States	www.ibm.com
InfluxData	2013	United States	https://www.influxdata.com/
ITRS	1993	United Kingdom	https://www.itrsgroup.com/
jKool	2014	United States	https://www.jkoolcloud.com/
Loom Systems	2015	United States	https://www.loomsystems.com/
Moogsoft	2011	United States	https://www.moogsoft.com/
Scalyr	2012	United States	https://www.scalyr.com/
ScienceLogic	2003	United States	https://sciencelogic.com/
SignalFx	2013	United States	https://signalfx.com/
Splunk	2004	United States	https://www.splunk.com/
Stackstate	2015	Netherlands	https://www.stackstate.com/
Sumo Logic	2010	United States	https://www.sumologic.com/
VNT Software	2010	Israel	http://www.vnt-software.com/
VuNet	2014	India	http://www.vunetsystems.com/

Source: Gartner (November 2018)

Market Recommendations

Ensure Success in the Deployment of AIOps Functionality by Adopting an Incremental Approach

The effective deployment of AIOps functionality requires a structured approach starting with the reorganization of IT domains according to data sources. This approach shifts the required focus toward datasets as opposed to tools, which has been the traditional approach. Gartner has found that it is best to begin with mastering the use of large persistent datasets ingested from a variety of sources. Only after the IT operations team has become fluent with the big data aspect of AIOps should it attempt mastery of the capability categories (see “12 Steps to Artificial Intelligence for IT Operations Excellence”). Hence, when selecting tools or services, an enterprise should prioritize those vendors that allow for the deployment of data ingestion, storage and access, independent from the remaining AIOps components. Given that AIOps will be used for multiple use cases, I&O leaders must ensure that the vendors support the gradual addition of those other functionalities.

Select AIOps Platforms Capable of Supporting a Broad Range of Historical and Streaming Data Types

Modern IT operations aim to gain a composite visibility to IT entities, including applications, their relationships, interdependencies and past transformations to gain insight to the present state of the overall IT landscape. Selection of the right data source is crucial in avoiding blind spots. The progressive nature of deployment maturity and evolving use cases requires a readiness to ingest a variety of data sources. I&O leaders must select AIOps platforms that are capable of ingesting and providing access to a broad range of historical and streaming data types (see Figure 3).

Figure 3. Data Types for Ingestion in AIOps

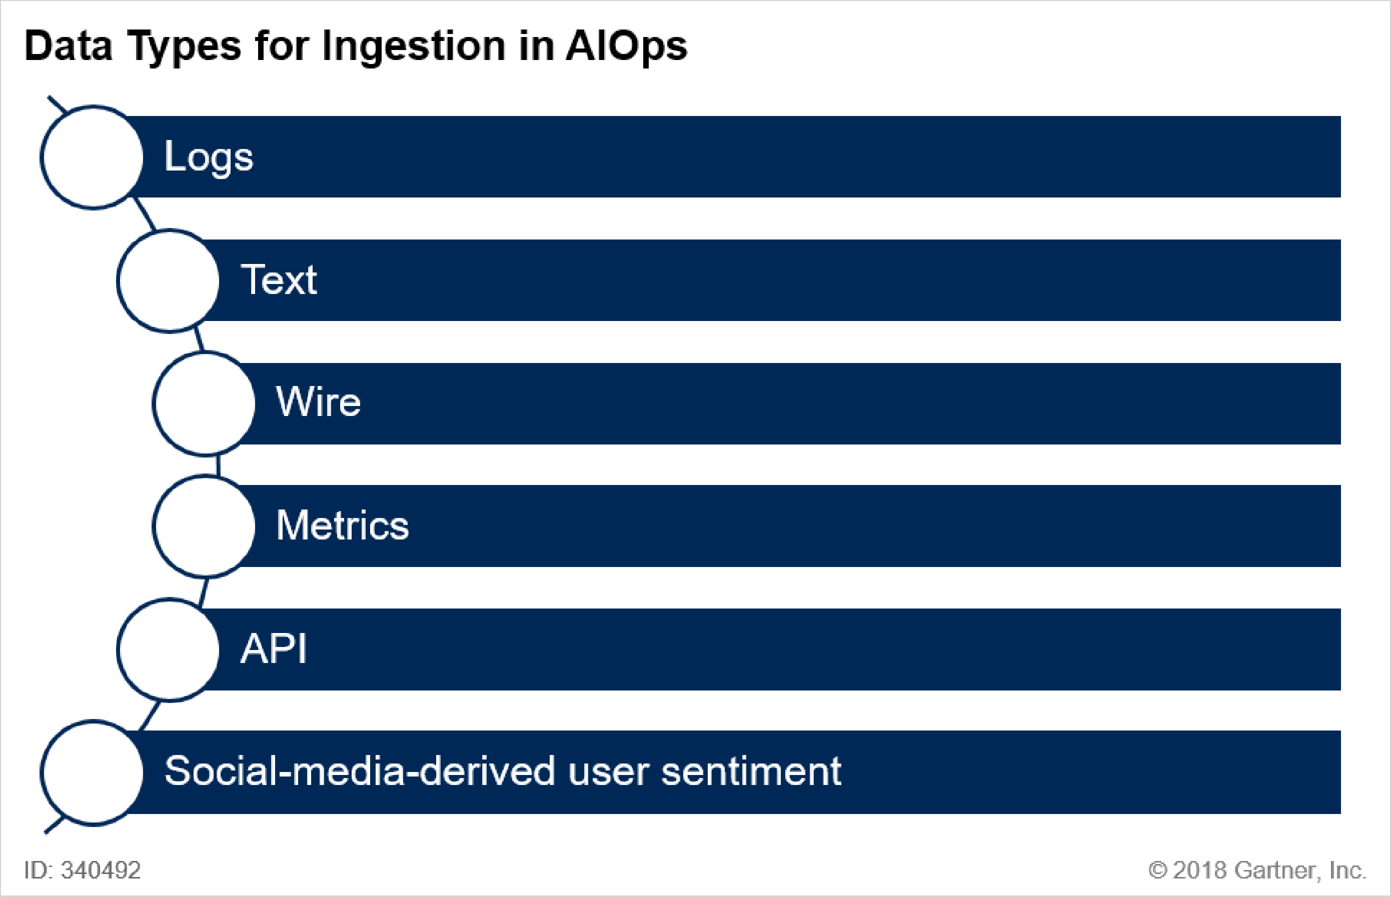
Data Types for Ingestion in AIOps
Logs
Text
Wire
Metrics
API
Social-media-derived user sentiment
ID: 340492
© 2018 Gartner, Inc.

Source: Gartner (November 2018)

AIOps platforms have historically focused on a single data source like logs or metrics. Unfortunately, no matter how large or frequently updated a given dataset is, restriction to a single data type tends to limit the insights into system behavior. Modern IT systems — with their modularity and dynamism — require a multiperspective approach to understand what is happening as they are being observed.

As an example, AIOps can be used to provide a consolidated analysis of digital experience across multiple digital services.

For example, use a pattern detection algorithm to improve the customer relationship process by detecting the patterns of behavior customers expressed in digital experience monitoring data. Use the machine learning algorithms in AIOps to perceive the patterns that relate user navigation with:

- Digital experience data from APM
- Order data pulled from payloads in business transactions
- Sentiment data from social media
- Service desk requests and statuses
- Account activity from the CRM system

Utilize this approach to build a composite model of a customer across all applications they use and even different behaviors across multiple modes of a single application (for example, when they use a web browser versus a mobile device). This can be used to predict customer churn and provide the insight and time necessary to prevent customer loss.

Choose Tools Offering the Ability to Systematically Progress Across the Four Phases of IT-Operations-Oriented Analytics and Machine Learning

Tools that support the incremental deployment of the four phases of IT-operations-oriented machine learning must be given a higher priority for investment (see Figure 5).

Figure 4. Four Phases of IT Operations

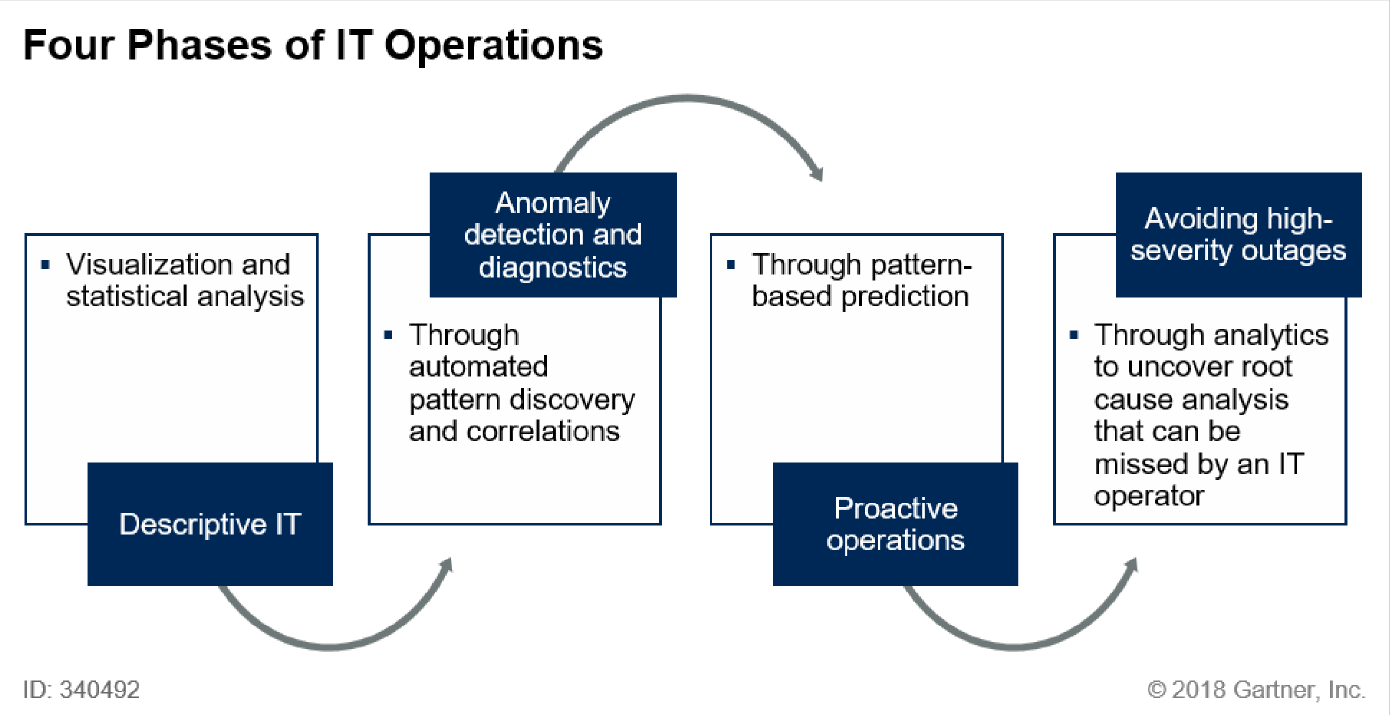
Four Phases of IT Operations
Visualization and statistical analysis
Descriptive IT
Anomaly detection and diagnostics
Through automated pattern discovery and correlations
Through pattern-based prediction
Proactive operations
Avoiding high-severity outages
Through analytics to uncover root cause analysis that can be missed by an IT operator
ID: 340492
© 2018 Gartner, Inc.

Source: Gartner (November 2018)

One of the key attributes toward enhancing IT operations teams' skills is an incremental approach.

The deployment of AI in an IT operations context is difficult and must be approached gradually. IT operations teams should begin their AI journey by becoming adept at data visualization and the use of basic statistical analysis. Resist the temptation to do it all at once. Only after these core "manual" disciplines have been mastered should machine learning proper be approached as follows:

- Initially experiment with allowing the software to reveal patterns that organize large volumes of data.
- Next, test the degree to which those patterns allow them to anticipate future events and incidents.
- Finally, work with root cause analysis functionality.

All four of these phases of AIOps are important, and enterprises should select tools that support as many as possible. These phases should be deployable in a modular manner, but also to ensure that IT operations can obtain value as they learn.

Evidence

[1] Over 400 inquiries over the past 12 months covering various aspects of IT monitoring and AIOps, including:

- Platform selection
- Deployment strategy
- Multiple AIOps use case within and outside IT to aid visualization, decisions and diagnostics

[2] 6% of the interactions related to AIOps were on various use cases for ITSM

[3] 3% of the AIOps interactions were related to the DevOps use case

[4] 15% of the interactions were related to the potential use of AIOps for customizing dashboards across various personas

Note 1
Representative Vendor Selection

The vendors listed in this research are picked as a sample based on one or two of the following criterion:

- Ability to ingest data from multiple sources, including historic and real-time streaming.
- Different offerings that include proprietary, open source, free and commercialized versions, including deployment that cuts across on-premises and SaaS-based options.

Note 2
Data Types

- **Log data ingestion —** Software that allows for the capture of alphanumeric text strings from log files generated by any software or hardware device, and the preparation of that data for access and analysis, indexed for storage.
- **Wire data ingestion —** Software that allows for the capture of packet data direct from taps on the network. All protocol and flow information should be prepared for access and analysis, and indexed for storage.
- **Metric data ingestion —** Software that allows for the direct capture of numerical data (for example, the capture of data to which time series and more general mathematical operations can be immediately applied).
- **Document text ingestion —** Software that allows for the ingestion, parsing, and syntactical and semantic indexing of human readable documents. This may include the use of technologies

commonly described as natural language processing (NLP).

About Careers Newsroom Policies Site Index IT Glossary Gartner Blog Network Contact Send Feedback

Gartner