GIGAOM

KEY CRITERIA/MARKET LANDSCAPE

Key Criteria for AIOps v1.0

DAVID S. LINTHICUM | MAY 14, 2020 - 4:55 PM CDT

TOPICS: ARTIFICIAL INTELLIGENCE DEV & OPS

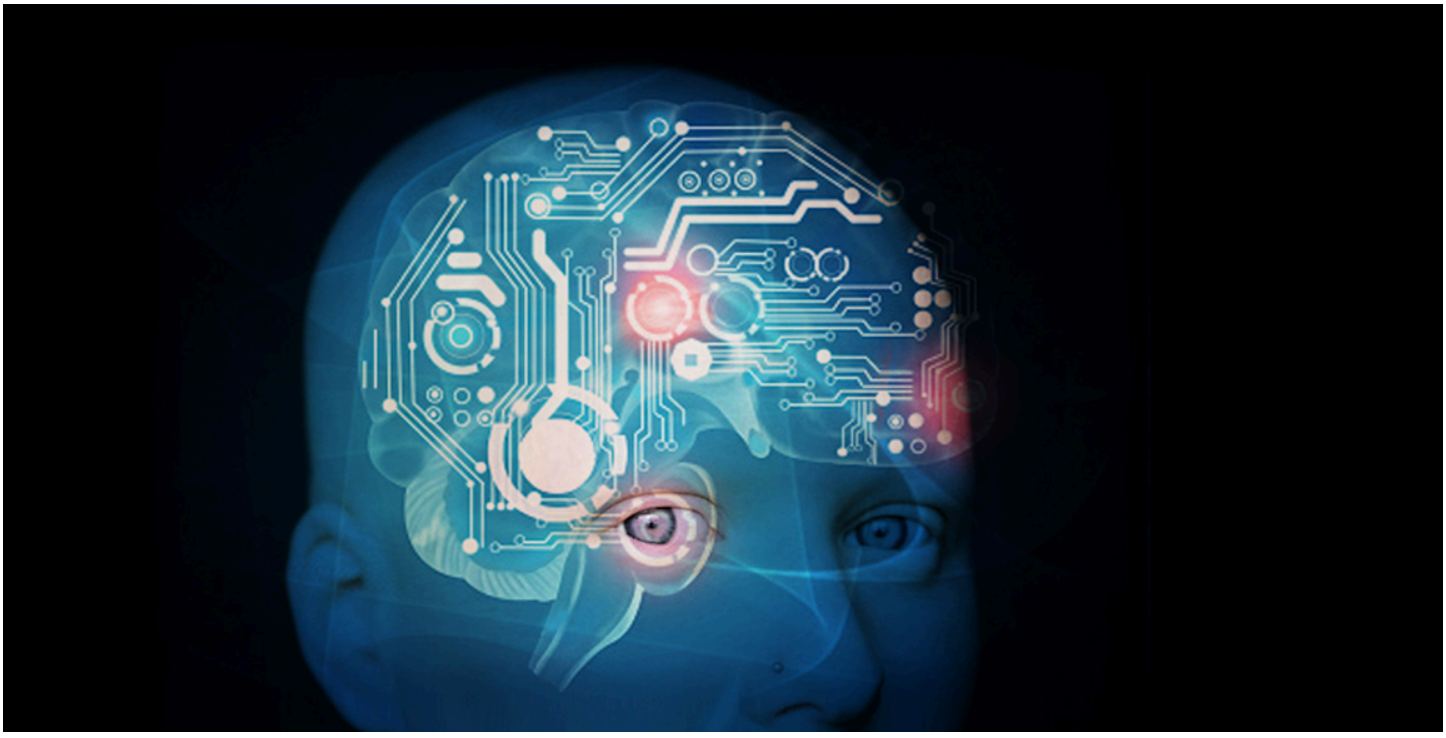

GIGAOM

Key Criteria for AIOps

TABLE OF CONTENTS

1 Summary

2 About the Key Criteria

3 Report Methodology

4 Primer: AIOps

5 Decision Criteria Analysis

6 Conclusion

7 About David Linthicum

8 About GigaOm

9 Copyright

1. Summary

AIOps tools, as the name implies, leverages a range of AI capabilities to enhance IT operations, including knowledge-driven predictive analytics and natural-language processing. AIOps is a technology that can be applied to all types of CloudOps tasks. These tools are built either to automate simple tasks so that IT operators can focus on more strategic work, or to perform tasks that are beyond a human's capabilities.

One can think of AIOps as a cross between active tools and those that can learn from being active. This is an important distinction. Tools must carry out pre-programmed, self-corrective processes, and it's the AIOps tools' ability to learn during these processes that creates a huge advantage. For instance, an AIOps tool might understand that performance issues could be the result of saturation caused by cyber-attacks, and that the situation should kick off security processes to mount a defense. For traditional tools, such an incident would be addressed as a simple performance issue, and not recognized as a security threat.

Moreover, and most important, AIOps tools have the ability to deal with thousands of data points and make correlations that most humans would not make. For instance, data update errors that lead to a pattern, and then leads to the identification of a bad network connection that would normally take weeks to diagnose.

However, the world of AIOps presents a duality. On the one hand, it's an emerging technology that for the first time mashes up operations and AI. On the other, many of the solutions in this space are traditional tools that have been updated to leverage AI. This mix of old and new, traditional players and startups, makes this space particularly interesting. Will traditional ops tools perhaps have more maturity and connections into specific systems, and therefore thrive? Or will new purpose-built tools fully leverage AI technology to enable more innovative approaches?

Conclusions reached in this report include:

- The AIOps tools in the market today are on a spectrum with regard to use of AI. While some make use of knowledge engines systemically in the monitoring and management of cloud and non-cloud systems, most tools leverage AI as an afterthought, not driving much of the functionality of the tool.
- Enterprises are typically adopting AIOps as an upgrade to existing ops tools, and are remaining brand loyal. This means that the upstarts in the AIOps space will find it difficult to break into a market where the established players are in essence selling with the same basic message: AI integrated with management and monitoring that you trust. Considering this, we may see a consolidation next year as the market focuses on a handful of players, down from the two dozen or so relevant players today.
- There seems to be two directions in AIOps: self-healing and not self-healing. Some AIOps systems are able to heal issues with systems that are managed and/or monitored. This means that if the tool finds an issue, a process is launched to attempt to correct the problem, for instance restarting a

server or a network hub. Other solutions are more passive, alerting users about an issue, but without taking automated corrective action. The trend is toward active, or self-healing, AIOps tools.

- These tools are all about the data. They store data as they monitor systems and can determine issues that need immediate attention, such as a down storage server. Or, they can deeply analyze historical data to determine trends that may portend a failure or other potential issue. The lifeblood of any AI system is the data needed to train the AI model, and this is the opportunity presented to AIOps tools. Monitored cloud or on-premises systems spin-off gigabytes of data each week, and that data can be fed into analytic systems augmented by AI.
- Enterprises that wish to leverage these tools should be careful to understand their capabilities, and should also test the tools across both enterprise cloud and non-cloud platforms. There have been compatibility issues reported, most discovered after deployment.
- Many of these tools are moving to an "on-demand" model, meaning that they will offer cloud-based services. This is an opportunity for those that have, or will have, the majority of their systems on public clouds. However, it may not be a good model for those that still have the majority of systems on-premises.

2. About the Key Criteria

HOW TO READ THIS REPORT

This GigaOm report is part of a series of documents that help IT professionals understand, explore, and evaluate a specific technology and its attendant market. It enables organizations to assess competing solutions in the context of well-defined criteria and metrics. For a fuller understanding, consider reviewing the following reports:

Key Criteria report: A detailed market sector analysis focused on a specific technology domain. The report enables IT decision-makers to make better decisions by defining key features and criteria for a product sector and assessing their impact on core evaluation metrics. This framework provides a strong overview of a technology sector and the solutions and vendors enabling it. The Key Criteria report is critical to informing the GigaOm Radar report.

Radar report: A market landscape analysis that provides a forward-looking evaluation of vendors and their solutions in a specific technology sector. The GigaOm Radar leverages scoring and qualitative analysis to plot a chart that depicts the relative value, character, and progression of vendors' solutions. The Radar report includes a breakdown of each vendor's offering in the sector.

Vendor Profile: An in-depth vendor analysis that provides an accessible, deep dive into a company's engagement with a technology sector. The analysis builds on coverage presented in the Key Criteria and Radar reports, drilling into details of the vendor's solution and assessing the company's strategy as it relates to the market sector. This analysis includes forward-looking guidance around both strategy and product.

3. Report Methodology

A Key Criteria report analyzes the most important features of a technology category to help the reader understand how they impact an enterprise and its IT organization. Features are grouped into three categories:

1. Table Stakes
2. Key Criteria
3. Emerging Technology

The goal is to help organizations assess capabilities and build a mid-to-long-term infrastructure strategy. In a mature technology, the solutions are divided into three target market categories: enterprise, high-performance, and specialized solutions. In a mature market, these differ in their characteristics and how they integrate with existing infrastructures. That said, the assessment is dependent more on the specific user's needs and not solely on the organization's vertical.

Table Stakes

Table stakes are system characteristics and features that are important when choosing the right solution. They include architectural choices that depend on the size of the organization, the requirements, the expected growth over time, and the types of workloads. Table stakes are mature, and the implementation of these features will not add any business advantage nor significantly change the Total Cost of Ownership (TCO) or Return on Investment (ROI) of the infrastructure.

Key Criteria

Key criteria features differentiate one solution from another. Depending on real user needs, they have a positive impact on one or more of the metrics mentioned. Therefore, implementation details are essential to understanding the benefits relative to the infrastructure, processes, or business. Following table stakes and key criteria, aspects like architectural design and implementation regain importance and need to be analyzed in great detail.

Emerging Technology

In this report section, we analyze exciting technologies on the horizon over the next 12 to 18 months. Some are already present in some form, but usually as part of niche products or suited for addressing very specific use cases. In either case, at this stage, the available implementations are not mature enough to be grouped in Key Criteria. Yet when implemented correctly and efficiently, this technology can make a difference to the metrics.

Over time, emerging technologies are refined and adopted to become key criteria, while key criteria

evolve and lose their status as a differentiator to become table stakes, as shown in **Figure 1**. Therefore, to get the best ROI, it's important to check up on what vendors are offering today and what they plan to release in the near future.

Figure 1: Evolution of Key Criteria and Emerging Technologies Over Time

4. Primer: AIOps

AIOps as both a concept and a tool category have evolved from the maturation of operational tools in general. Over the past few years, most providers in the traditional ops tools space have bolted AI engines onto their tools and called it AIOps, no matter if they leverage AI systemically or not. However, there are some purpose-built AIOps tool startups out there that are leveraging AI from the jump.

All Ops tools have been data gathering and analytics tools from the beginning, thus adding AI to traditional or new Ops tools allows them to now learn from that data. In some cases, the tools can correct issues using pre-programmed routines, such as restarting a server or blocking an IP address that seems to be attacking one of your servers.

This provides a few advantages:

1. We can remove the humans from CloudOps processes for the most part, only alerting when things require human intervention. This means fewer operational personnel and lower costs.
2. We can integrate AIOps with other enterprise tools, such as DevOps or governance and security operations.
3. We can look for trends that allow the operational team to be proactive. For example, the AIOps tool can monitor a networking switch that's about to fail and is putting an increasing amount of errors on the network.

Now that we're a few years into this paradigm and its technology offerings, we're starting to note some patterns—some good, and some not so good:

- AIOps tools in many instances are ops tools in their fourth, fifth, or sixth generations, and this maturity means that many of the bugs have been worked out of the basic monitoring subsystem.
- Most of the AIOps tools have had public cloud management in mind for a while and are able to bridge the gap between on-premises legacy system management and managing applications and services in the public clouds.
- The AIOps tools are capable tools for managing and monitoring cloud, multi-cloud, legacy, and even IoT and edge-based systems.
- The ability to support complex system heterogeneity is really the true value of the ops tools, and why they are important to those implementing cloud or non cloud systems.

Of course, there are some downsides:

- Most AIOps users are not taking advantage of the AI subsystems in the tools, so you may be paying for a feature you're not using.

- Lack of training, or a lack of valid use cases in the current set of systems—cloud and not cloud. In other words, there may be no need for AI at all, and it's "shelfware" for now.

AIOps Categorization

AIOps is really an existing category of tools, CloudOps and Ops tools, repurposed with the AI subsystems. This is leading to a number of new abilities, such as:

- **Predictively spot when traditional systems or cloud-based components are about to fail:** This is achieved by looking at the patterns of data of similar servers by leveraging machine learning, which led up to a failure in the past. The idea is that we can go beyond the abilities of traditional operational tools leveraging AI, now providing the ability to dynamically predict issues, in many cases before they cause operational outages.
- **Self-Healing:** Meaning that upon spotting the issues with the cloud-based or on-premises component, the tool is able to take pre-preprogrammed corrective action, such as restarting a server or disconnecting from a bad network device. This relieves the humans of dealing with the issues ongoing, and should address 80 percent of the ops tasks, now automated for all but the most critical issues.
- **Connecting to remote components, inside and outside of public clouds:** This is all about gathering and analyzing data in real-time, keeping historical data paths, and learning from that data at the same time by leveraging a machine learning engine. Thus, the ability to connect into remote components, such as servers and networking devices, is critical to an AIOps tool being effective.
- **Creating customized views to promote productivity:** Information views should be configurable for specific roles and tasks. The DBA, for instance, should be able to see the database servers under AIOps management, and see very detailed information about the health of those servers, such as memory usage, cache usage, and production of data. However, others on the Ops team may find that more general views are more productive, where they can observe the systems holistically.
- **Monitoring, managing, and repairing standard infrastructure concepts:** This means the ability to gather operational data from storage, network, compute, data, applications, and security systems.

While it is difficult to pin down exactly what is meant by an AIOps tool, patterns are beginning to emerge that allow those watching the market to place them into specific categories, sometimes more than one. For our purposes we can divide AIOps into four categories: Active, Passive, Homogeneous, and Heterogeneous (See **Figure 2**).

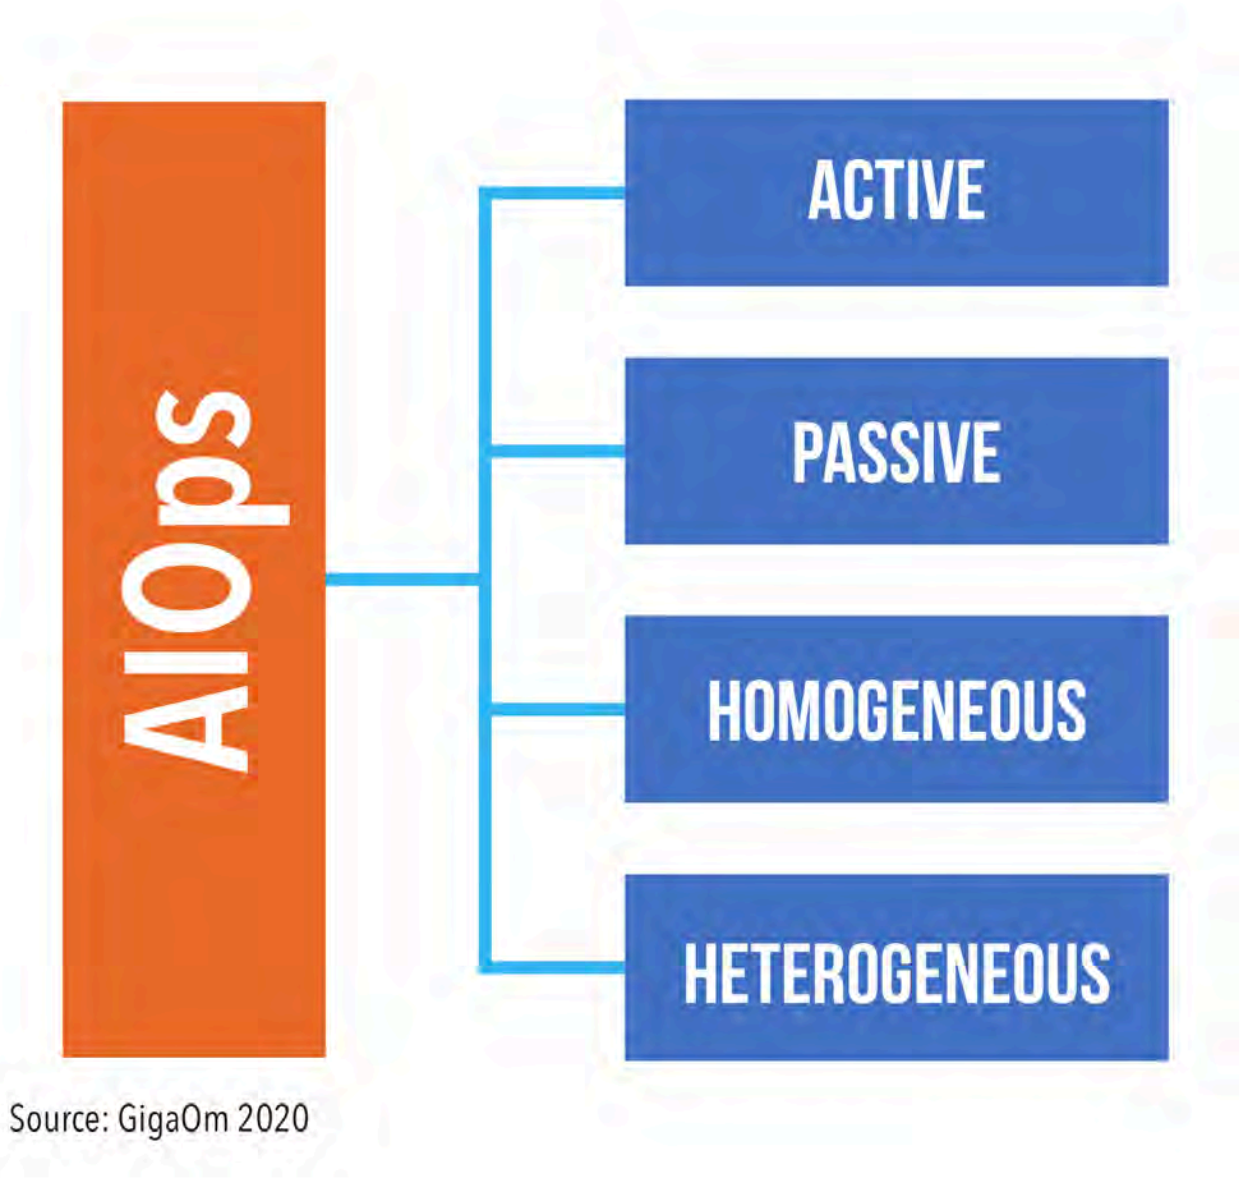

Figure 2: Four AIOps Categories

Active

Active refers to tools that are able to self-heal system issues discovered while monitoring. This means that they can understand that an issue exists, either through trending or spotting data points that are out of a predetermined threshold. For example, the database server is producing corrupted data. Once the AIOps tool determines this, it can trigger a set of automated scripts to potentially fix the issue, such as reindexing the database, resetting the database, or even failing over to the replicant database existing in another cloud region.

This proactive automation of correcting ops issues is really where the value of AIOps exists, considering that it allows enterprises to hire fewer ops engineers, and increase up-time significantly. While most enterprises don't consider these as value drivers when selecting a AIOps tool, most active AIOps tools should pay for themselves within a few months.

Passive

Passive AIOps tools, as you may have guessed, don't provide the ability to take corrective action. While many of these AIOps providers establish partnerships with tool providers that do, enabling that functionality then becomes a DIY proposition.

Passive AIOps tools are largely data oriented, and spend their time gathering information from as many data points as they can connect to. They also provide real-time and analytics-based data analysis to enable impressive dashboards for operational professions. However, despite the ability to do alerting, they can't carry out corrective action natively.

Homogeneous

As the name implies these are AIOps tools that live on a single platform. For example a native Ops tool that employs AI native to a single cloud provider. While the tool can manage services, such as storage, data, compute, and the like that are native to the cloud brand, the inability to span different platforms for operational management imposes obvious limitations for those servicing a hybrid or multi-cloud deployment.

Heterogeneous

Of course most of the AIOps tools are heterogeneous, meaning that they are able to monitor and manage a variety of different cloud brands, as well as native systems operating within the cloud providers. Moreover, and perhaps of more value, these AIOps tools can manage traditional on-premises systems and even mainframes, as well as IoT and edge-based computing environments.

AIOps Break Down

As you can see in **Figure 3,** AIOps tools have several core components. The repository that stores configurations, policies, and other information around what's connected and how they are managed. Also, there is typically some sort of database that stores the data coming in from the connected systems. Keep in mind that while some tools work directly from the downloads, or linking to the SaaS service, most of them require custom configuration, which can range from very easy, to very time-consuming. You need to consider these issues during the tool evaluations.

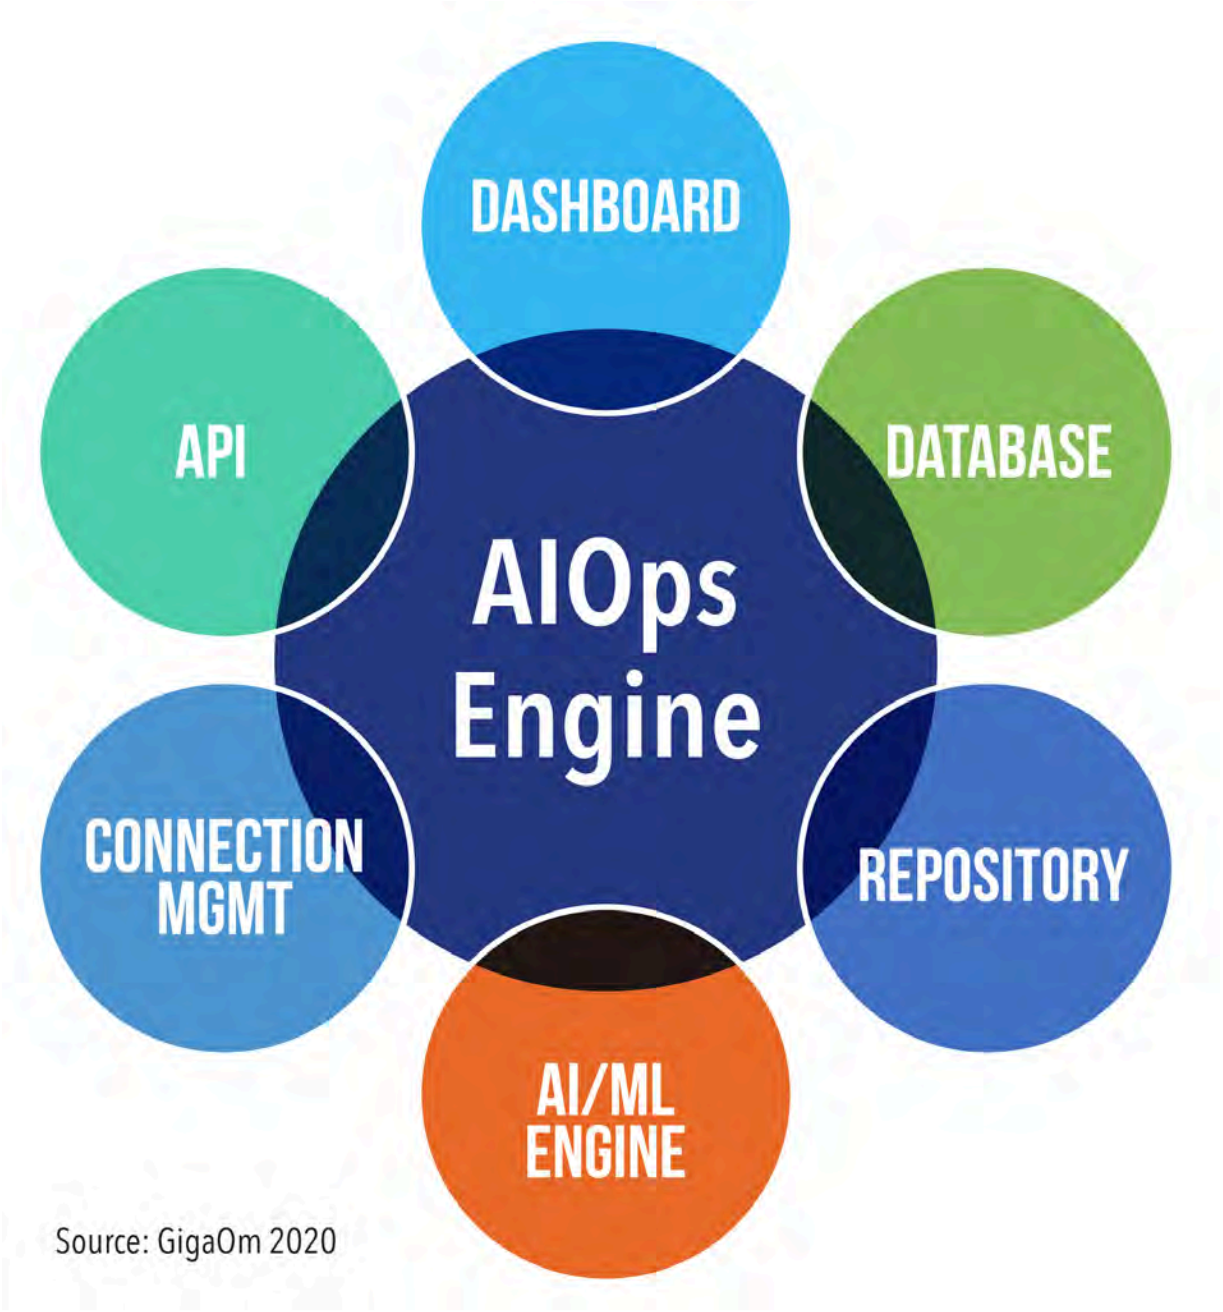

Figure 3: Components of an AIOps Tool

The dashboard is where the operational analytics are presented, which is typically customizable to meet the specific needs of the Ops teams. An API system enables external applications to leverage the services and data contained within the AIOps system. Also, a connection manager handles the connectivity between the systems that produce data and the database that stores it. Finally, but most importantly, the AI engine that's available to be trained by the incoming data, and provides services internal to the AIOps tool.

5. Decision Criteria Analysis

In this section, we describe in detail the criteria that organizations should use to evaluate AIOps solutions. First, we describe table stakes criteria common to most solutions, before focusing on the key criteria that are most impactful in differentiating and selecting a solution. We also describe emerging technologies and features. Finally, we explore evaluation metrics--broad product characteristics that reflect fundamental aspects of a product or services, such as scalability or cost-effectiveness.

Table Stakes

While there are many different features that AIOps tools provide, it's helpful to find common patterns and features that allow you to break down the capabilities for better AIOps evaluations and selection driven by business requirements.

AIOps tools provide five common patterns, or table stakes, as presented in **Figure 4**.

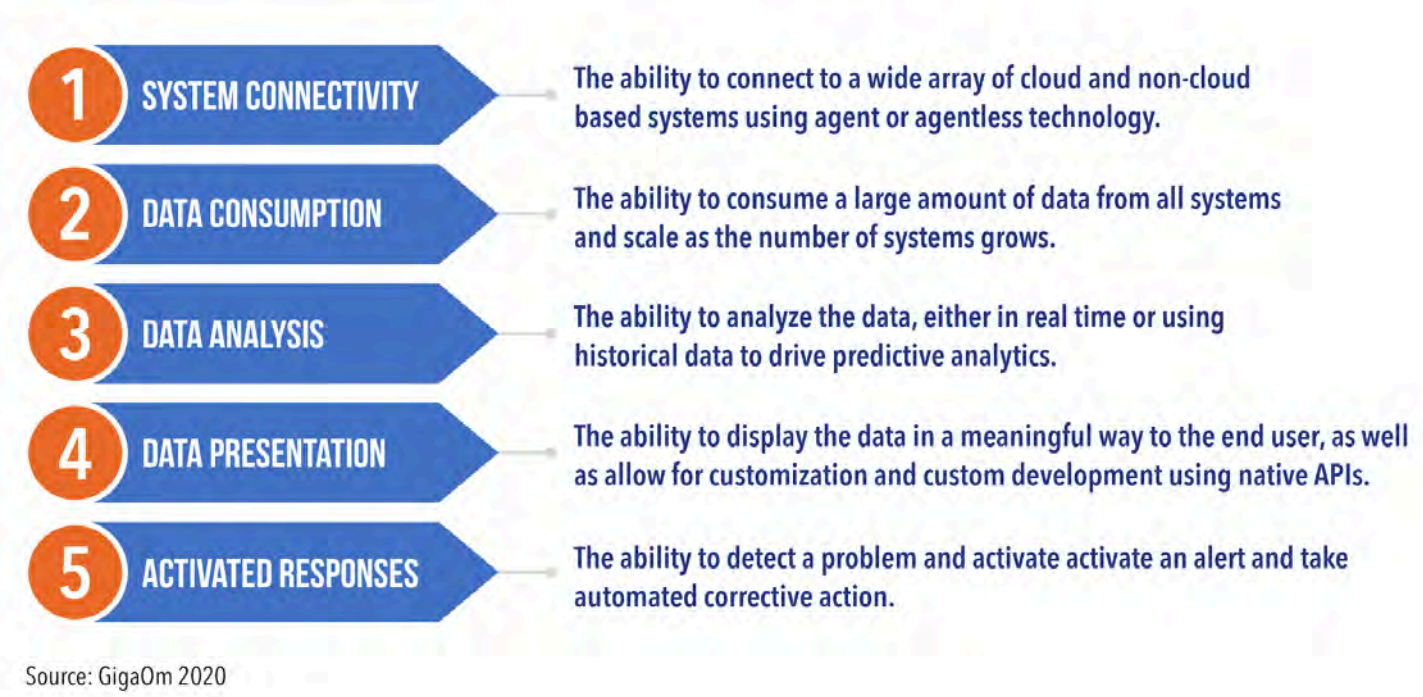

Figure 4: AIOps Table Stakes

System Connectivity

This is often overlooked in terms of what systems the AIOps tool can monitor. If you're running cloud or on-premises systems such as storage, compute, databases, or even applications, and the solution does not support those systems directly, you will be paying for custom development to make the connection. This is not advisable.

Data Consumption

The ability to consume data being sent from the source systems under management. This can be all types of data, but must be placed in a structure where it can be analysed either by the AIOps tool itself, or by external analytics tools. Moreover, this is where the AI engine can learn from the data and leverage it during operations.

Data Analysis

Along with the AI engine, data analysis means that the tool can make good use of the data it's gathered from the systems under management. For instance, the ability to determine if a DDOS attack is underway, providing alerts as well as the ability to defend the system automatically. The AI engine is typically coupled to this feature.

Data Presentation

This is the subsystem that's able to present the data to the Ops team member, and do so in such a way that it has meaning to the team member. This means the ability to customize the dashboards, such as only providing database operational data and analytics to the database operations person, with other data points not appearing and thus distracting from their focus.

Activated Responses

This is the ability to carry out some pre-programmed action as a part making an operational conclusion that something is in alarm. Response can vary from texting the admin to alert him or her of an operational issue, to invoking a script that attempts a repair or recovery automatically.

Key Criteria

Key criteria for selection of this technology include:

Proactive/Self-Healing Operations

The AIOps tool is able to solve problems automatically, and without human intervention. While some

tools can do this trick natively, others are only able to do this with integration of another tool. An example of proactive/self-healing operation would be the ability to reset a virtual storage service when I/O begins to fail.

Learning Systems

The AI engine is able to learn from the data being consumed by the AIOps tool. The key feature here is the ability to react to learned data, such as spotting a server that's about to go down based on similar historical patterns. Learning systems range from very advanced to loosely integrated, and need to be tested before the tools are accepted.

A related concept, Continuous Improvement of tools**,** determines if an AIOps tool is able to consistently improve its function. For example, automatic updates that include enhancements and fixes, such as updating the AIOps tool's configuration using automated methods, or even updating training data to allow the AI engine to make better decisions. Finally, the progression of a centralized knowledge model that will be leveraged by a brand of an AIOps tool, that should provide shared knowledge bases that all AIOps tool users can leverage.

Configurability of Dashboards

The dashboard is customizable and provides the ability to leverage advanced analytics tools, including creation and maintenance of algorithms and policies. Instances of these include the ability to create separate dashboards for mobile devices or desktop computers.

Heterogeneous Monitoring

The AIOps tool is able to monitor and serve compute environments that span both on-premises and cloud. This is important as existing on-premises systems increasingly work with public cloud-based systems.

System Connectivity Support (Agent and Non-Agent)

Enable connectivity to a wide variety of systems, such as storage, compute, and applications, and ensure that those connections are maintained. Self healing is important here, considering that connections are often lost and need to be recovered automatically.

Integration with Enterprise Systems

The AIOps tool can share data and services with other tools and applications in the enterprise. This is typically accomplished through a well-defined API, but can also be carried out using native connectivity with the AIOps tool.

Component Monitoring

In addition, there are three key criteria that are bound together under the rubric of component monitoring. They are:

Cost and usage monitoring: The ability to monitor usage and cost. While typically focused on public clouds and usage-based billing, this system can be applied to traditional systems as well. The result should be a holistic look at what things cost, and who's doing what.

End-user monitoring: The tool can monitor and manage what the end user is experiencing, and may provide some self-healing capabilities, such as the ability to prompt the end user to launch a corrective action on demand. This functionality must adjust to the different platforms that the end users leverage, including mobile devices and even IoT devices.

Application monitoring: While the AIOps tool can see the infrastructure resources, there is also visibility directly into the applications. This means that application-level management can occur via an interaction with the AIOps system.

Emerging Technologies

If there are any game-changing aspects of AIOps they include:

Centralized Knowledge Sharing

The AIOps tool can benefit not only from data and learning it collects from your organization over time, but from the data and learning shared by other enterprises as well. This has the ability to revolutionize the use of AIOps, considering you won't have to wait months for the AIOps tool to get smart by processing locally generated data.

Tight Integration with Security and Governance

Considering that Ops and security should work together, there is no greater advantage than to have an Ops tool that can talk to the security manager and the other way around. This will provide the advantage of leveraging operations as a first line of defense, such as shutting down a server that's under attack, or shutting off access to a storage system that's compromised.

Extending to Mobile Computing and Beyond

While mobile is a whole other management and operations story, we can expect AIOps tools to extend not only to mobile computing platforms, but to other emerging platforms such as IoT and edge computing.

Evaluation Metrics

Having explored AIOps features and criteria across table stakes, key criteria and emerging technologies, we turn our attention to evaluation metrics. These are overarching characteristics of a solution that are impacted and defined by the criteria we've documented. For instance, an AIOps solution that supports both learning systems and continuous improvement of tools might be judged highly on the metric of flexibility, while one that is limited in heterogeneous monitoring and integration with other systems may do poorly on the metric of interoperability.

While each evaluation carried out by an enterprise will vary due to different requirements, there are some general characteristics appropriate to AIOps solutions to consider. These include:

- Flexibility
- Scalability
- Interoperability
- TCO/ROI
- User Experience
- Ease of Use
- Ecosystems

For purposes of this report, and the related GigaOm Radar Report for AIOps, it is important for decision makers to consider the following specific evaluation metrics when assessing AIOps solutions. These include:

Number of Systems Supported

This is simply the number and types of systems a solution provides connectivity with out of the box. Also to be considered is the type of connectivity: Agent-based, meaning that a small piece of software runs on the system being monitored, or agentless, meaning that there is no agent required. System connectivity support, integration with enterprise systems, and heterogeneous monitoring are key criteria that impact this evaluation metric.

Management Approaches

This describes how a vendor approaches operations from the unmodified tool. This information should include most operational patterns and best practices. The key criteria to consider here include configurability of dashboards and component monitoring.

Learning Approaches

Considering that we're leveraging a native AI system as part of AIOps, we must also understand how the AIOps tool learns over time, and the processes that are leveraged to do so. Among key criteria presented in this report, learning systems has the most direct impact on this evaluation factor.

Overall Operational Impact

This describes how much ROI the AIOps tool is expected to drive over time. This metric is typically measured in avoided downtime, based on the fact that the AIOps tool is going to be more proactive in detecting and resolving error and failure states. Virtually every key criteria listed in this report will impact the operational impact and ROI produced by an AIOps solution, however, proactive/self-healing operation and the nature of the learning system may stand apart in enabling advanced functionality.

6. Conclusion

While AIOps as a category is fairly new, the concept of operational tools certainly is not. Indeed, IT organizations have been using Ops tools (albeit, in traditional ways) for decades. Now, however, vendors are leveraging AI capabilities in a variety of ways to provide a more effective approach and technology set to deal with the rising complexity of sophisticated distributed systems, such as multi-cloud deployments.

In order to take advantage of this technology for successful operations, you must arm yourself with the knowledge presented in this report, which should enable you to select the right tool, or at least to determine if these tools are equipped to solve your issues.

In truth, there is no magic bullet here. These are simply tools that can leverage insight gleaned from data that's been collected over time. In other words, the longer you leverage an AIOps tool, the more valuable it should become as the knowledge engine learns through experience, much as we do in our jobs.

7. About David Linthicum

David Linthicum is a CTO and internationally renowned thought leader in cloud computing. David has spent the last 25 years leading, showing, and teaching large global enterprise organizations across all industries how to use technology resources more productively and constantly innovate.

David has been a CTO five times for both public and private companies, and a CEO two times in the last 25 years. David has published 13 books on computing and his thought leadership has appeared in Wall Street Journal, NPR, Forbes, InfoWorld and Lynda.com. He has expanded the vision of both startups and established corporations as to what is possible and achievable.

All of David's opinions are his own.

8. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

9. Copyright