GIGAOM

MARKET RADAR

GigaOm Radar for Cloud Observability v1.0

ANDY THURAI | FEB 26, 2021 - 10:57 AM CST

TOPIC: CLOUD INFRASTRUCTURE

CREDIT: BAGOTAJ

GigaOm Radar for Cloud Observability

TABLE OF CONTENTS

1. Summary

Observability is an emerging set of practices, platforms, and tools that goes beyond monitoring to provide insight into the internal state of systems by analyzing external outputs. It's a concept that has its roots in 19th century control theory concepts and is rapidly gaining traction today.

Of course, monitoring has been a core function of IT for decades, but old approaches have become inadequate for a variety of reasons—cloud deployments, agile development methodology, continuous deployments, and new DevOps practices among them. These have changed the way systems, infrastructure, and applications need to be observed so events and incidents can be acted upon quickly.

At the heart of the observability concept is a very basic premise: quickly learn what happens within your IT to avoid extended outages. And in the unfortunate event of an outage, you need to ensure that you can get to the root cause of it fast. Outages are measured by Mean Time To Resolution (MTTR) and it is the goal of the observability concept to drive the MTTR value to as close to zero as possible.

No surprise, building resilient service delivery systems that are available with high uptime is the ultimate end goal for any business. Achieving this goal requires executing three core concepts:

- **Monitoring:** This is about understanding if things are working properly in a service-centric manner.
- **Observability:** This is about enabling complete end-to-end visibility into your applications, systems, APIs, microservices, network, infrastructure, and more.
- **AIOps:** This is about using comprehensive visibility to derive meaning from the collected data to yield actionable insights and courses of action.

To achieve observability, you need to measure the golden telemetry signals—logs, metrics, and traces. Logs and metrics have been measured by IT professionals for decades, but traces is a fairly new concept that emerged as modern applications increasingly were built using distributed microservices. A service request is no longer completed by one service but rather by a composition of microservices, and as such there is an imperative to track or trace the service request from start to finish. In order to generate proper telemetry, all the underlying systems must be properly instrumented. This way enterprises can achieve full visibility into their systems to track service calls, identify outages, and determine if the impacted systems are on-premises, in the cloud, or somewhere else.

Observability is not always about introducing new tools, but about consolidating the telemetry data, properly instrumenting systems to get the appropriate telemetry, creating actionable insights, and avoiding extended outages. Comprehensive observability is core to future proofing IT infrastructure.

This report evaluates key vendors in the emerging application/system/infrastructure observability space and aims to equip IT decision makers with the information they need to select providers according to their specific needs. We analyze the vendors on a set of key criteria and evaluation metrics, which are described in depth in the "Key Criteria Report for Cloud Observability."

HOW TO READ THIS REPORT

This GigaOm report is one of a series of documents that helps IT organizations assess competing solutions in the context of well-defined features and criteria. For a fuller understanding consider reviewing the following reports:

Key Criteria report: A detailed market sector analysis that assesses the impact that key product features and criteria have on top-line solution characteristics—such as scalability, performance, and TCO—that drive purchase decisions.

GigaOm Radar report: A forward-looking analysis that plots the relative value and progression of vendor solutions along multiple axes based on strategy and execution. The Radar report includes a breakdown of each vendor's offering in the sector.

Vendor Profile: An in-depth vendor analysis that builds on the framework developed in the Key Criteria and Radar reports to assess a company's engagement within a technology sector. This analysis includes forward-looking guidance around both strategy and product.

2. Market Categories and Deployment Types

For a better understanding of the market and vendor positioning (**Table 1**), we assess how well solutions for cloud observability are positioned to serve specific market segments.

Deployment Type

Observability tools are generally delivered in the following deployment and consumption models:

Public SaaS
The platform is managed by the vendors on public cloud locations (AWS, GCP, or Azure) and offered as a SaaS offering to users. It can be accessed directly through a web portal with no additional installation. This is often the simplest and easiest way to leverage observability and is the defacto standard for observability platform deployment by most vendors.

Private SaaS
A variation of the public SaaS model, private SaaS vendors offer their observability solutions as hosted, multi-tenant SaaS platforms from their own data centers. While this approach offers most of what is available in public SaaS, some cloud-native capabilities may be lacking—for example, support for serverless and the availability of additional tools in the public cloud.

On-Premises Software
Some vendors offer an option to install, configure, and self-manage the solution on private cloud locations (such as on-premises). This can be a compelling option for customers concerned about security and compliance, for example, but can also impose restrictions on usage flexibility and may lack cloud-native capabilities.

Instrumentation and Openness

Openness
Some observability vendors offer either free open-source or free and open-license solutions with significant add-on capabilities. A good example of the latter approach is the solution offered by Elastic built on the ELK (Elasticsearch, Logstash, Kibana) stack. Others provide a combination of open-source integrations to craft best-of-breed implementations and make them enterprise grade by offering enhanced security, scalability, manageability, governance, and other capabilities. An example of this approach would be Logz.io, which is built on the ELK stack and incorporates Prometheus and Grafana along with Jaeger for distributed tracing.

Instrumentation
This is an area where there is a lot of variation. Some vendors require installation of instrumentation to collect telemetry. Most times this process is manual and can be overwhelming especially when installed on distributed microservices. Some vendors offer one-click or auto instrumentation features that can make this process easier. Some vendors also offer integration with open source tools such as

FluentD and Prometheus. Some of these choices can be very limiting and others quite open and flexible.

OpenTelemetry

The OpenTelemetry Initiative provides an open-source observability framework for cloud-native software that is expanding rapidly to include open standards for logs, metrics, and traces. Big-name cloud players are moving to embrace OpenTelemetry quickly and observability vendors are likewise offering integration with OpenTelemetry tools.

The level of integration varies from vendor to vendor, with some offering only data exchange formats to support OpenTelemetry standards while others offer complete open source/open telemetry integration. Full adoption of the OpenTelemetry standards can yield significant benefits around instrumentation, as customers can deploy drop-in instrumentation regardless of the platform. Portability becomes achievable as well, improving both cost savings and efficiency. With OpenCensus and OpenTracing merging, and Jaeger tracing integration gaining support, the most difficult part of observability—tracing—is coming together well.

Table 1: Vendor Positioning

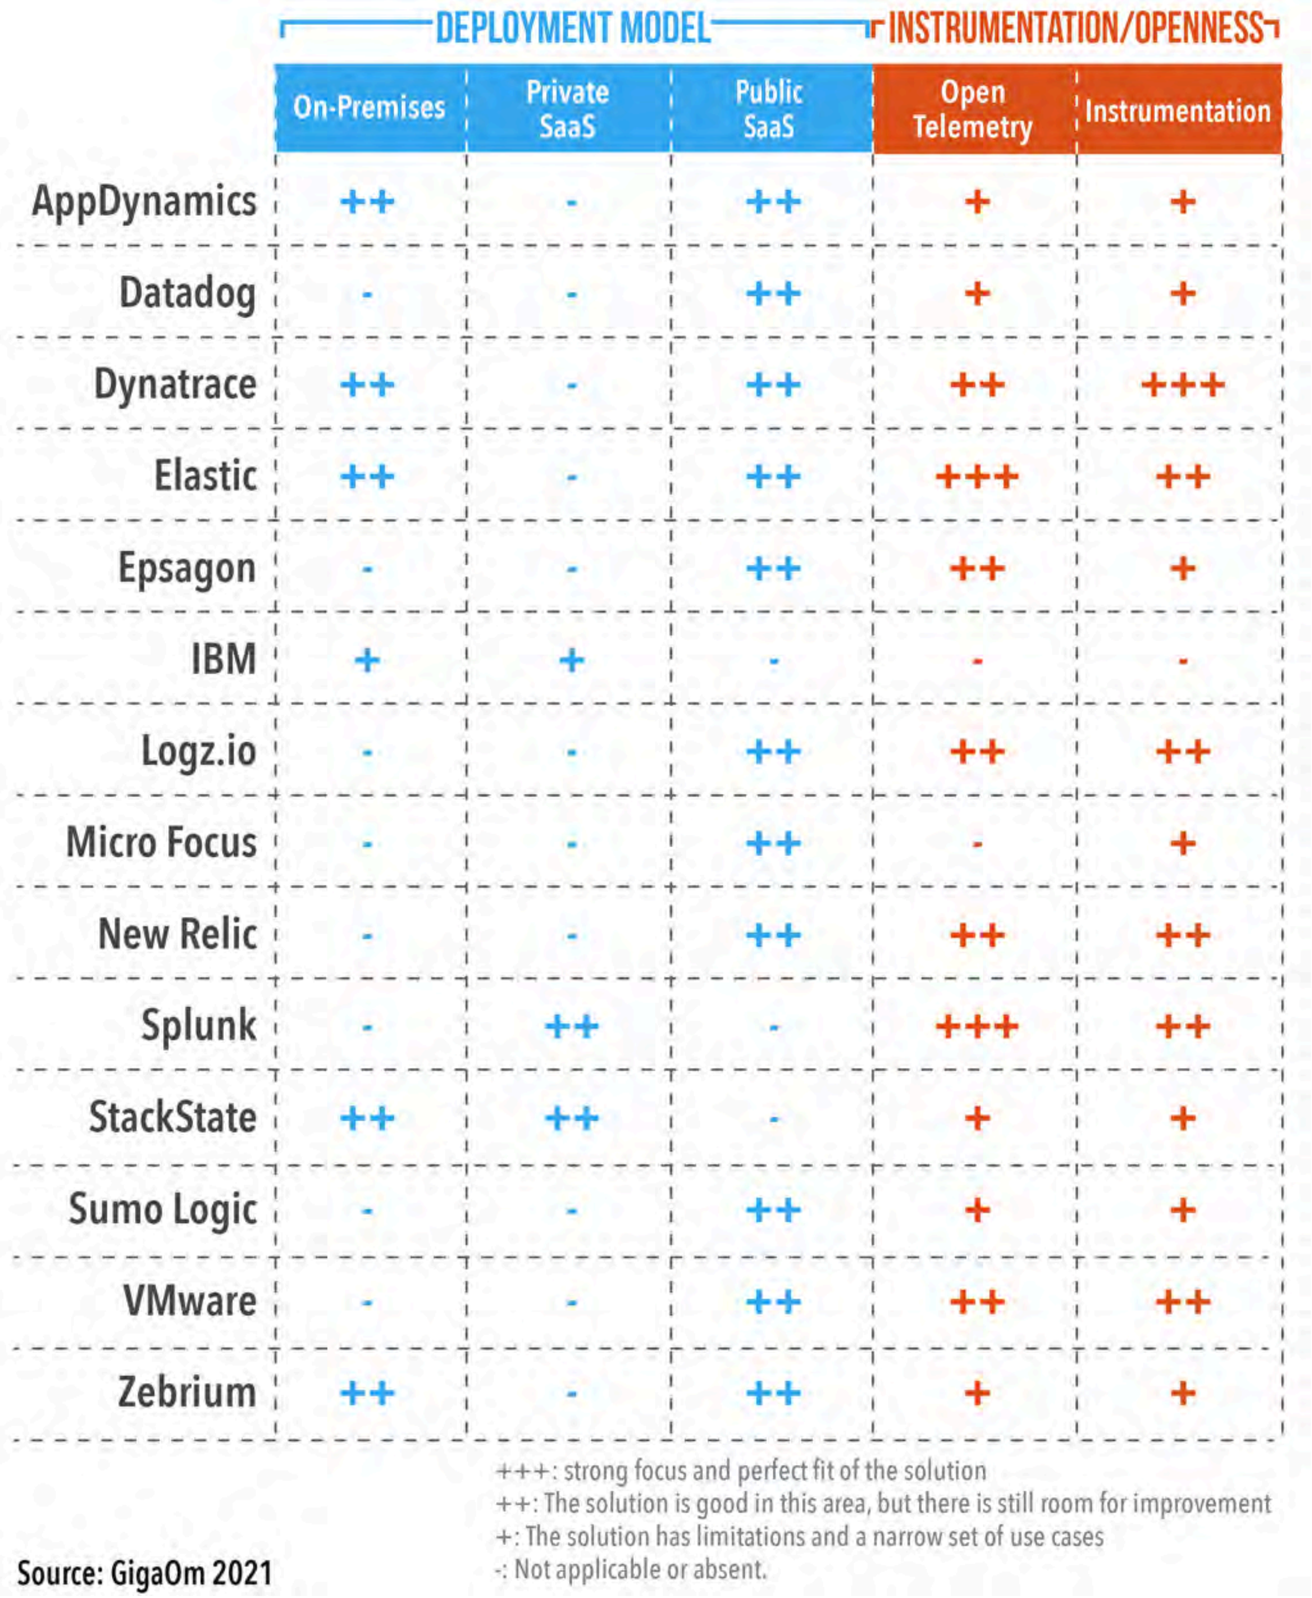

	DEPLOYMENT MODEL			INSTRUMENTATION/OPENNESS	
	On-Premises	Private SaaS	Public SaaS	Open Telemetry	Instrumentation
AppDynamics	++	-	++	+	+
Datadog	-	-	++	+	+
Dynatrace	++	-	++	++	+++
Elastic	++	-	++	+++	++
Epsagon	-	-	++	++	+
IBM	+	+	-	-	-
Logz.io	-	-	++	++	++
Micro Focus	-	-	++	-	+
New Relic	-	-	++	++	++
Splunk	-	++	-	+++	++
StackState	++	++	-	+	+
Sumo Logic	-	-	++	+	+
VMware	-	-	++	++	++
Zebrium	++	-	++	+	+

+++: strong focus and perfect fit of the solution
++: The solution is good in this area, but there is still room for improvement
+: The solution has limitations and a narrow set of use cases
-: Not applicable or absent.

Source: GigaOm 2021

3. Key Criteria Comparison

Building on the findings from the GigaOm report, "Key Criteria for Evaluating Cloud Observability," **Table 2** summarizes how each vendor included in this research performs in the areas that we consider differentiating and critical in this sector. The objective is to give the reader a snapshot of the technical capabilities of different solutions and define the perimeter of the market landscape.

Observability and resiliency have been material concerns in the IT space for decades, however the emergence of Google SRE best practices for building resilient systems helped bring these concepts to the forefront. With the introduction of Service Level Objectives (SLOs), the idea of cost-to-upkeep versus fixing technical debt is an interesting and relevant debate.

While evaluating vendors, we kept the dynamics of technical debt in mind as we assessed the differentiation of old technologies versus new technologies. We also gave strong consideration to vendors that offer open telemetry, open connectivity, open interfaces, and charge based on consumption models. The entire software world is moving toward pay-as-you-go, consumption-based pricing models, and the benefits this can provide in the observability realm are significant.

Table 2. Key Criteria Comparison

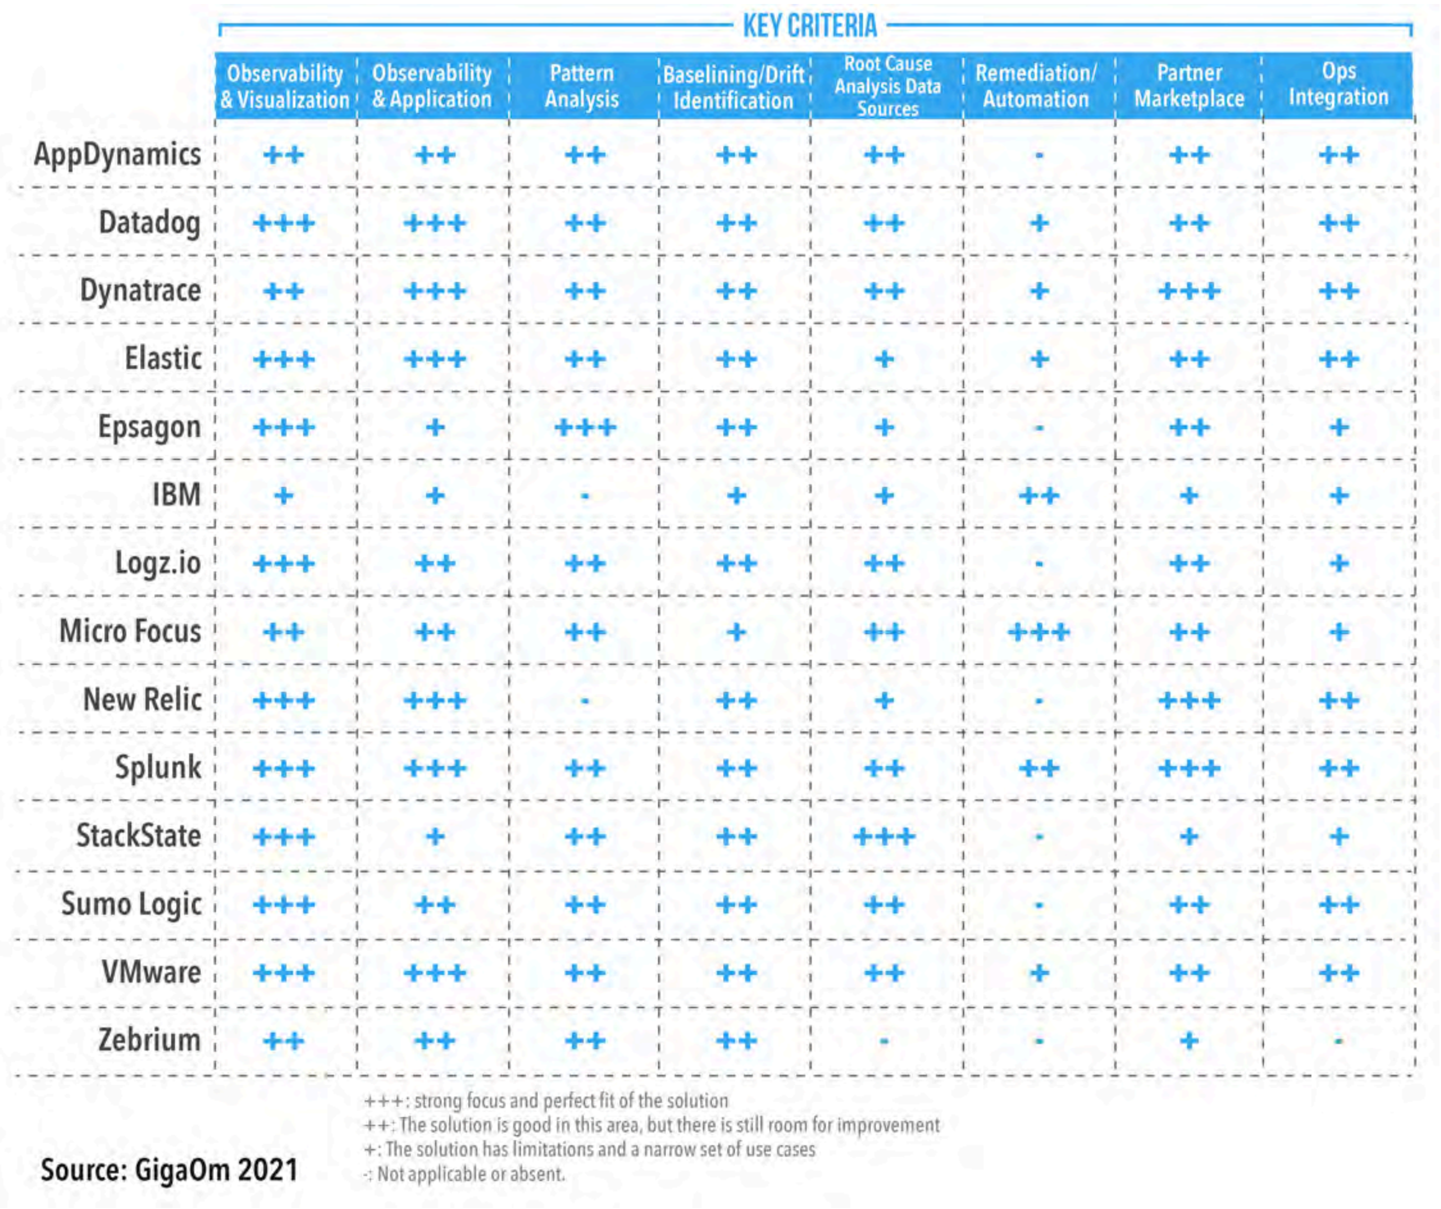

	KEY CRITERIA							
	Observability & Visualization	Observability & Application	Pattern Analysis	Baselining/Drift Identification	Root Cause Analysis Data Sources	Remediation/ Automation	Partner Marketplace	Ops Integration
AppDynamics	++	++	++	++	++	-	++	++
Datadog	+++	+++	++	++	++	+	++	++
Dynatrace	++	+++	++	++	++	+	+++	++
Elastic	+++	+++	++	++	+	+	++	++
Epsagon	+++	+	+++	++	+	-	++	+
IBM	+	+	-	+	+	++	+	+
Logz.io	+++	++	++	++	++	-	++	+
Micro Focus	++	++	++	+	++	+++	++	+
New Relic	+++	+++	-	++	+	-	+++	++
Splunk	+++	+++	++	++	++	++	+++	++
StackState	+++	+	++	++	+++	-	+	+
Sumo Logic	+++	++	++	++	++	-	++	++
VMware	+++	+++	++	++	++	+	++	++
Zebrium	++	++	++	++	-	-	+	-

+++: strong focus and perfect fit of the solution
++: The solution is good in this area, but there is still room for improvement
+: The solution has limitations and a narrow set of use cases
-: Not applicable or absent.

Source: GigaOm 2021

Table 3. Evaluation Metrics Comparison

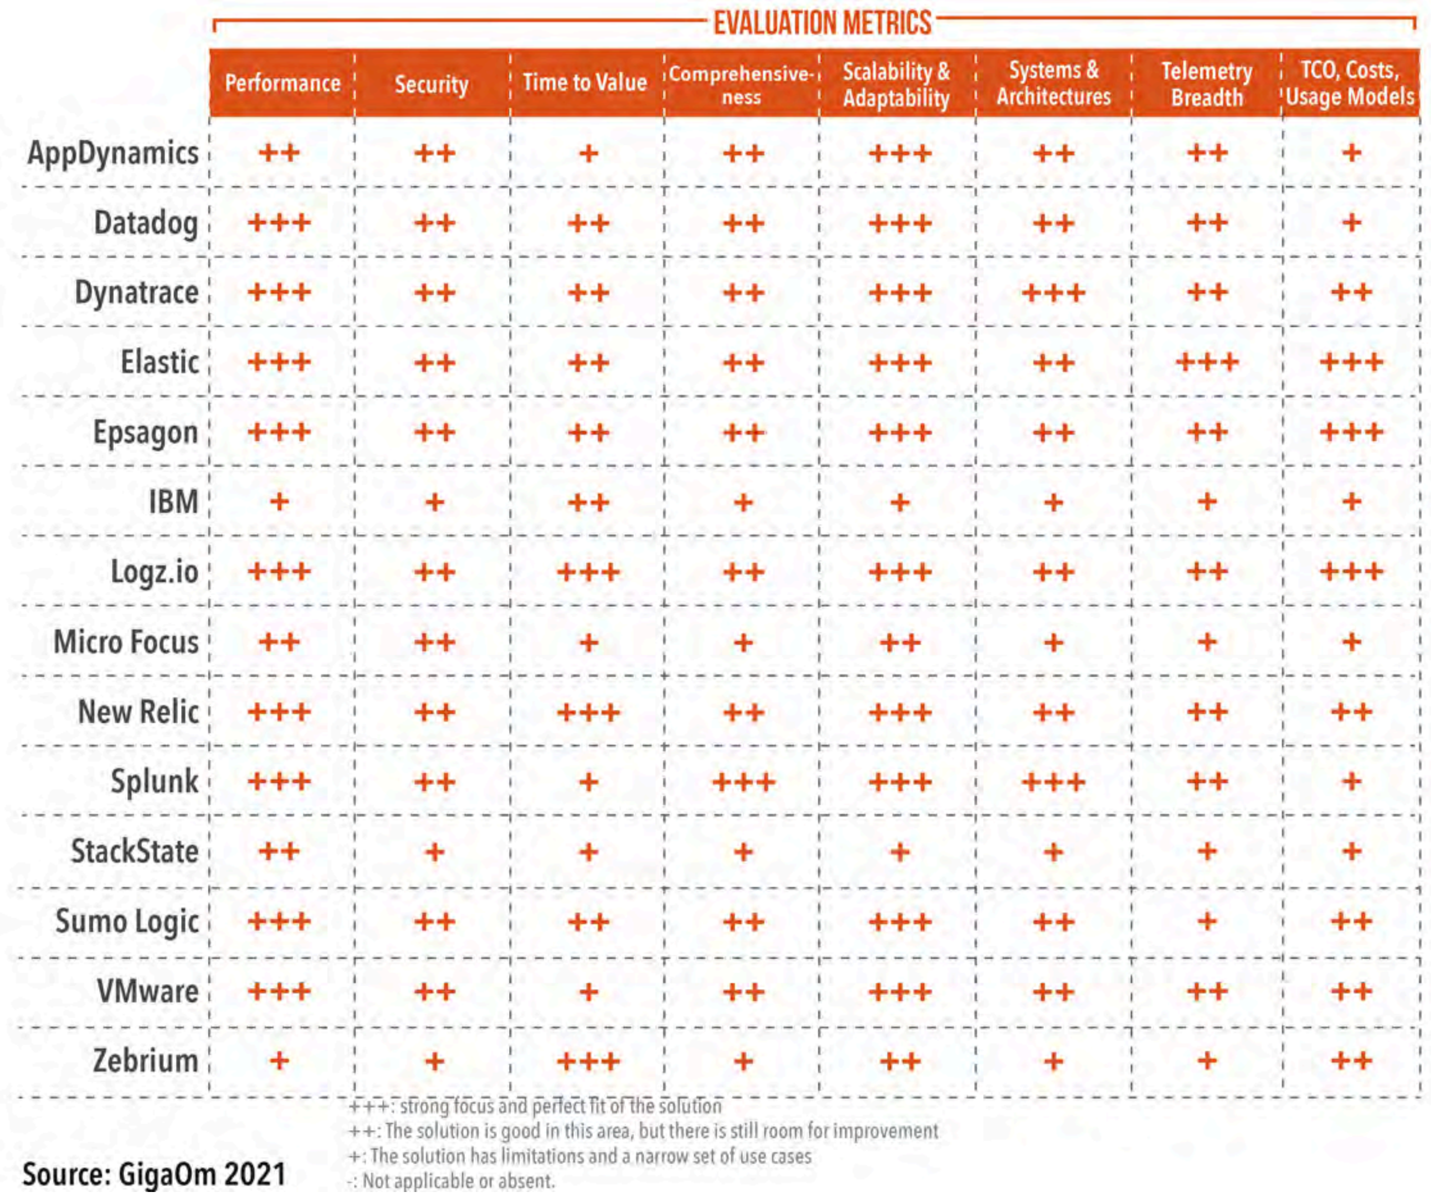

	EVALUATION METRICS							
	Performance	Security	Time to Value	Comprehensiveness	Scalability & Adaptability	Systems & Architectures	Telemetry Breadth	TCO, Costs, Usage Models
AppDynamics	++	++	+	++	+++	++	++	+
Datadog	+++	++	++	++	+++	++	++	+
Dynatrace	+++	++	++	++	+++	+++	++	++
Elastic	+++	++	++	++	+++	++	+++	+++
Epsagon	+++	++	++	++	+++	++	++	+++
IBM	+	+	++	+	+	+	+	+
Logz.io	+++	++	+++	++	+++	++	++	+++
Micro Focus	++	++	+	+	++	+	+	+
New Relic	+++	++	+++	++	+++	++	++	++
Splunk	+++	++	+	+++	+++	+++	++	+
StackState	++	+	+	+	+	+	+	+
Sumo Logic	+++	++	++	++	+++	++	+	++
VMware	+++	++	+	++	+++	++	++	++
Zebrium	+	+	+++	+	++	+	+	++

+++: strong focus and perfect fit of the solution
++: The solution is good in this area, but there is still room for improvement
+: The solution has limitations and a narrow set of use cases
-: Not applicable or absent.

Source: GigaOm 2021

By combining the information provided in the tables above, the reader can develop a clear understanding of the technical solutions available in the market.

4. GigaOm Radar

This report synthesizes the analysis of key criteria and their impact on evaluation metrics to inform the GigaOm Radar graphic in **Figure 1**. The resulting chart is a forward-looking perspective on all the vendors in this report, based on their products' technical capabilities and feature sets.

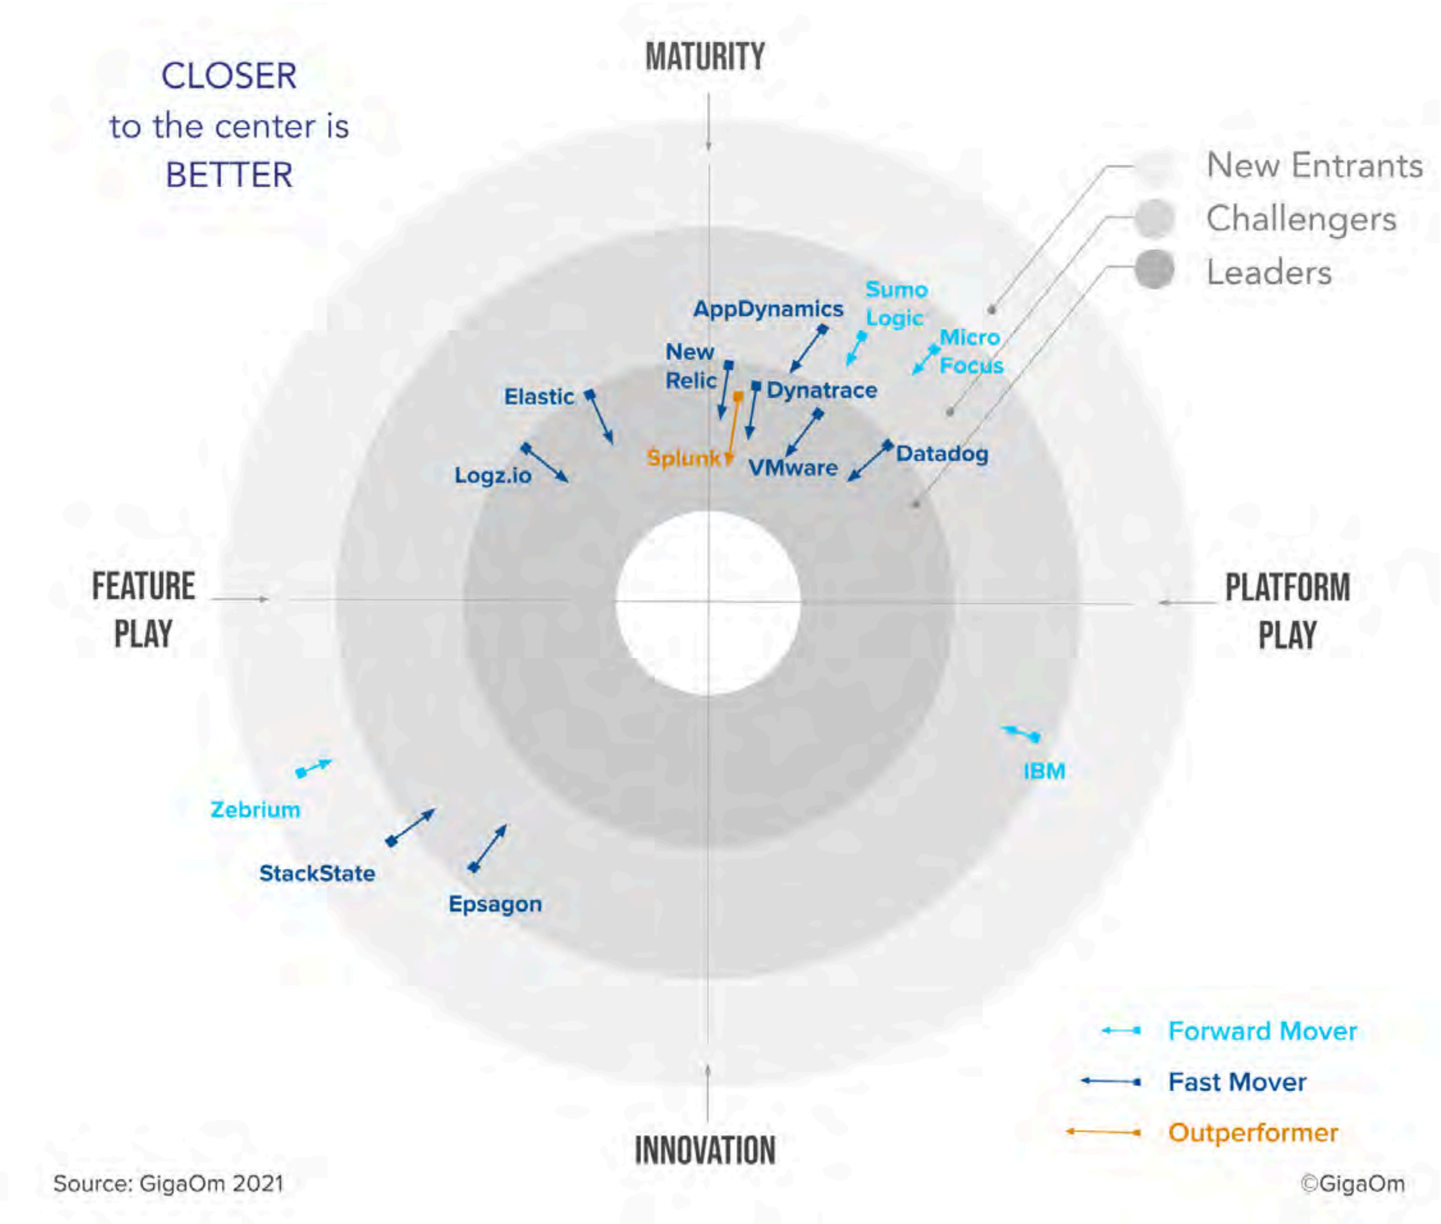

Figure 1: GigaOm Radar for Cloud Observability

The GigaOm Radar plots vendor solutions across a series of concentric rings, with those set closer to center judged to be of higher overall value. The chart characterizes each vendor on two axes--Maturity versus Innovation, and Feature Play versus Platform Play--while providing an arrow that projects each solution's evolution over the coming 12 to 18 months.

As you can see in the Radar chart in **Figure 1**, there is no dominant player in this report; rather a number of companies have established themselves as Leaders as they continue to add capabilities via acquisition, organic development, and partnerships. APM vendors have made strides by adding log and infrastructure metrics capabilities, while log vendors have pushed to add support for metrics and trace capabilities.

It is refreshing to see vendors expanding their capabilities to embrace cloud-based approaches. While legacy vendors may struggle to enable cloud-native technologies, they retain their value in providing enterprise functionality, integration, and robustness. What is clear is that no single vendor can provide a one-size-fits-all, magic bullet solution to observability. But a combination of solutions, or a vendor solution augmented with open source tools, can get you close—especially for AWS shops that have the option of committing to innovative AWS-focused solutions that offer small form factors and cost-effective approaches.

Our analysis shows that organizations must determine their priorities across use cases, location choices, functions and features, and finally cost to identify the optimal vendor and solution. Decision makers must be alert to changes in the market. Mergers and acquisitions are red hot in the observability space and it's producing some volatility. For instance, while we were able to take into account the Splunk/SignalFx merger, the acquisition of Instana by IBM occurred after we were briefed by the vendor. We've strived to capture the observability market as it exists today, while taking into consideration the forward-looking impact posed by vendor roadmaps.

Another factor worth mentioning is that most vendors tend to excel in certain areas and lag in others. For a vendor to be considered a Leader in this report, we valued consistency across logs, metrics, and traces over the entire stack, at both cloud and on-premises locations, while emphasizing support for open telemetry and ease of instrumentation at reasonable cost. This is a challenging blend of capabilities for vendors to execute.

INSIDE THE GIGAOM RADAR

The GigaOm Radar weighs each vendor's execution, roadmap, and ability to innovate to plot solutions along two axes, each set as opposing pairs. On the Y axis, **Maturity** recognizes solution stability, strength of ecosystem, and a conservative stance, while **Innovation** highlights technical innovation and a more aggressive approach. On the X axis, **Feature Play** connotes a narrow focus on niche or cutting-edge functionality, while **Platform Play** displays a broader platform focus and commitment to a comprehensive feature set.

The closer to center a solution sits, the better its execution and value, with top performers occupying the inner Leaders circle. The centermost circle is almost always empty, reserved for highly mature and consolidated markets that lack space for further innovation.

The GigaOm Radar offers a forward-looking assessment, plotting the current and projected position of each solution over a 12- to 18-month window. Arrows indicate travel based on strategy and pace of innovation, with vendors designated as Forward Movers, Fast Movers, or Outperformers based on their rate of progression.

Note that the Radar excludes vendor market share as a metric. The focus is on forward-looking analysis that emphasizes the value of innovation and differentiation over incumbent market position.

5. Vendor Insights

Cisco AppDynamics

With its acquisition of AppDynamics and ThousandEyes, Cisco has moved into the observability space. The AppDynamics solution is geared toward mid-sized to large enterprises and appeals to financial, retail, and IT services customers. The company has transitioned AppDynamics from an APM solution to an observability platform, adding capabilities that include cloud, network, and infrastructure monitoring. The solution can visualize revenue paths and correlate customer and application experience to find and fix app issues. It can also monitor errors using its cognition engine, isolate problematic domains, and identify root causes from snapshot data by scanning all instances of collected telemetry in the dependency tree using the Automated Transaction Diagnostic feature.

The APM feature affords visibility down to the code level and into important transactions across multi-cloud environments. The infrastructure monitoring tool provides users a view of connections between applications and infrastructure, whether the application is hybrid cloud, multi-cloud, or on-premises.

Cisco AppDynamics has the ability to ingest data from its own agents as well as via open standards such as Prometheus and OpenTelemetry. It also supports public clouds such as AWS, consumes up to 450 billion metrics a day, and can handle structured and unstructured data. Its systems do not use sampling.

Like other APM providers, the platform uses a topology and dependency-aware data model that spans domains. It can baseline any collected or computed metrics to learn what normal looks like. The AI automatically detects anomalies and removes them from the baseline computation to reduce false positives, normalizes reference data for higher quality, and identifies seasonality based on historical data to reduce manual intervention.

Cisco AppDynamics has bi-directional integration with ServiceNow, which maps CMDB configuration items with AppDynamics entities. This enables the visualization of application and supporting infrastructure health in ServiceNow and a fast, contextual navigation between platforms to speed troubleshooting.

Cisco AppDynamics also integrates with Cisco Intersight Workload Optimizer, with CI/CD platforms like Harness.io, and has a large array of APIs, allowing customers to integrate a wide range of other tools and platforms.

AppDynamics best suits mid-sized to large organizations that can exploit its features and AI engine. The solution is available as a SaaS/multi-tenant solution, but can be set up on-premises as well. Currently the pricing is set per-agent, which could be expensive for large operations, though Cisco is working on a new pricing model.

Traditionally an enterprise APM solution, Cisco has added functionality to observe modern applications as well. While AI/ML has been integrated into the platform, use cases around AIOps remain limited. The solution provides a decent set of visibility features for AWS applications but sorely lacks these for Azure and GCP. The company also needs to expand OpenTelemetry integration, open source collectors, and the like.

Strengths: Cisco acquired ThousandEyes, which adds synthetic monitoring for Internet and cloud environments and enhances AppDynamics digital experience monitoring. The visual dashboard experience offers service outage insight down to the application level, as well as insight into affected transactional costs. Transaction snapshotting lets developers perform deep code inspection of API calls.

Challenges: Given Cisco's dominance in the network field, it's surprising that the AppDynamics solution doesn't accept and co-relate more network data. Areas that could be improved include agent handling, integration of metrics with Azure, GCP, and other cloud providers, and log management (though the acquisition of dashbase.io should help here).

Datadog

Datadog was formed in 2010 with the aim of removing friction between developers and system administrators. Its growth is driven by a focus on automation and real-time observability. Launched as an infrastructure monitoring company, Datadog has expanded its portfolio via both acquisition and organic innovation to offer solutions in the full observability space.

Datadog monitors infrastructure systems using agent and API calls to support containers, VMs, services, databases, storage systems, and network devices. If an agent cannot be deployed, it can gather information remotely using SNMP, JMX, OpenMetrics, or remote API calls.

The Datadog agent runs inside customers' Kubernetes clusters to collect metrics, traces, and logs in real time. The agent can sit directly on the operating system, as a containerized deployment, sidecar deployment, or agentless using a Lambda layer in serverless environments.

The cloud-native network monitoring tool counts packets between source/destination IP/port endpoints. Using tagging, customers track dependencies and metrics such as request volume, RTT variance between persistent components like services or applications, and short-lived components like Kubernetes pods.

The platform offers more than 400 integrations, such as AWS Fargate and Lambda, Google Cloud Run and Functions, Azure Functions, and Azure App Service for serverless. It also offers orchestrators such as Kubernetes, OpenShift, Amazon ECS and AWS Fargate, Rancher, Mesos, Docker Swarm, Cloud Foundry, and Azure Container Instances.

The Datadog platform has a common unified, dynamic tagging structure across the platform, allowing it to triangulate and support RCA, unification, and telemetry across the stack and across the platform.

Datadog's Log Management strategy is built around Logging without Limits™, which gives customers flexibility and cost controls via prioritization of logs. APM and distributed tracing is done via Tracing without Limits™, which reduces the tradeoff between cost and visibility by allowing customers to send 100% of their traces, search and analyze them all in a 15-minute live rolling window, and then index the most important ones.

The platform automatically retains error and high-latency traces for 15 days. Traces can be searched and analyzed using tags and facets, and they are represented by the Datadog flamegraph that visualizes the execution path of the request. The Datadog application monitoring suite also includes Real User Monitoring (RUM) and Synthetics for front-end visibility.

Datadog offers a variety of pricing structures for solutions within its observability platform, be it infrastructure monitoring, log management, APM, or security monitoring. Datadog's setup is involved with its agent-based solutions. The solution can be deployed as a SaaS, multi-tenant solution. No on-premises versions are offered.

Datadog has significantly bolstered its original host-based infrastructure monitoring system. The additions of application insights, application/transaction tracing, log for associated traces, code profiling, real user/experience monitoring, network monitoring, and compliance monitoring are all welcome. However, the solution needs to mature and extend functionality to emerge as a first-class full stack enterprise observability solution.

Strengths: Full-fidelity data collection (down to one-second interval) and storage improves data analysis. Comprehensive AWS integration and decent support for on-premises and private cloud makes this a good combination for hybrid teams. Single agent to collect all telemetry eliminates the need to deploy a new agent every time new data needs to be collected.

Challenges: Solution needs to improve and mature, especially in application monitoring and tracing. Limited log management and search around events are concerns, as is lagging support for unstructured data needs. Pricing is somewhat high.

Dynatrace

Dynatrace has built a solid reputation as a high-quality application performance management (APM) solution. It is now building on that reputation with its full observability platform, which is based on Davis—the company's proprietary AI engine.

The Dynatrace platform includes APM, AIOps, infrastructure monitoring, digital business analytics, and digital experience management (DEM) for enterprise IT departments and digital businesses. Using automation in concert with the Davis AI engine, the Dynatrace platform provides root-cause details of application performance, generates insights into the underlying infrastructure, and presents an overview of the user experience. The system is designed to scale and perform on-premises or in hybrid, cloud, or edge environments.

The platform's OneAgent drops a single binary onto a host to automatically instrument not only containers running within the environment, but also processes running within the container; all without requiring any manual instrumentation or image modification.

Auto-adaptive AI/ML models trigger 90 percent of all detected and analyzed problems, along with their graph-based root-causation process. Davis can evaluate how serious an issue is based on user impact, transaction value, lost productivity, and other factors.

Dynatrace supports end-to-end tracing for all queue technologies based on JMS. It also supports IBM MQ, Apache Kafka, RabbitMQ, ActiveMQ, MSMQ, Tibco EMS, and Tibco Rendezvous.

The Dynatrace solution is best suited for medium-to-large enterprises that need to automate observability processes. It is available as a full SaaS or distributed SaaS solution. The distributed SaaS version requires a private cluster, which can be on-premises or in a public or private cloud. The full SaaS platform receives updates every two weeks, and the distributed version every month.

The Dynatrace cluster is typically set up on a single node and can be extended to multiple nodes as needed. High-availability capabilities and load balancing are built-in. The solution creates a DNS name and certificate automatically for a secure out-of-the-box experience.

A pioneer in the APM space, Dynatrace re-architected its solution a few years ago and has since emerged as a strong player in the observability sector. Integration with more than 500 out-of-the-box solutions allows for easier connections within the ITOps space. Full data collection (without sampling) is welcome, though it can produce enough metrics and tracing data to potentially overwhelm existing systems.

Dynatrace supports OpenTelemetry and is one of it's top five contributors. While Dynatrace is not a logging solution per se, it does integrate well with common log providers. The mainframe-to-microservice topology mapping and business impact dashboards make the solution stronger.

Strengths: Using a single agent to instrument an environment is an interesting concept. There is currently full open telemetry-native integration for metrics, and that is coming soon for traces. The dashboard feature is particularly cool, and demonstrates business, application, and user impact while replaying transactions in real time.

Challenges: The solution is node-based. It's not architected on Kubernetes or any other container, but can run natively in the cloud. It also lacks integration with other strong pure-play AIOps players, which could be useful for many customers.

Elastic

Elastic has built a solid observability platform using the free and open ELK Stack (Elasticsearch, Logstash, Kibana). The company has successfully layered usability and visibility on top of the stack and its technology is used widely across enterprises as diverse as eBay, Wikipedia, Uber, and Netflix.

Elastic offers both enterprise and cloud (AWS, Azure, and GCP) versions. This helps users create independent, hybrid-cloud, or multi-cloud variations of the solution as needed. This is particularly useful when an enterprise needs to start at one location (either on-premises or in the cloud) and quickly expand to other locations without creating siloed implementations or fragmenting the toolset.

Elastic Observability uses its own popular open source ingestion products, Logstash and Beats, for data collection and forwarding. Beats are lightweight shippers that collect and ship data from edge location, and Logstash is a server-side data processing pipeline that supports multiple sources and sinks. This provides a strong combination for search, combined observability, endpoint security, and SIEM

The solution provides visibility across the entire ecosystem, using the underlying Elasticsearch to quickly search for relevant insights within all data. The solution combines logs, metrics, synthetics, and traces in the Elasticsearch datastore. It also uses Kibana for quick results, reducing MTTR and improving “war room” collaboration.

Elastic Observability aims to create a centralized knowledge resource, making it easier to view and investigate all aspects of complex infrastructure. It supports open standards and agents such as OpenTelemetry and Jaeger, as well as APIs to increase portability.

Elastic uses the Elastic Common Schema to normalize data and make seach, analysis, error detection, and alert functions quicker and simpler. Kibana provides clear visualizations to facilitate analysis. Machine learning ensures clear alerts and provides opportunities for automation.

Elastic's unsupervised machine learning model takes a baseline of observed systems and identifies behavioral anomalies. Users can control baseline frequency according to their needs (per service, per host, per environment, and so on), but the default setting is slightly less than 10 minutes.

Elastic Observability enables remediation automation with the most common ITSM and DevOps workflow automation vendors (such as ServiceNow). It also uses manual analyst-assisted remediation (via PagerDuty, Slack, and Jira), alert webhooks for generic remediation integrations (including use of custom scripts), and API-level integration for the most customized automated remediation approaches.

The popularity of the Elastic Stack is largely driven by its ability to scale without impacting performance, via the scale-out architecture of each layer. Elastic Observability retains that strength. It can scale up and down with no service interruption.

Network monitoring is also part of the solution. While not as comprehensive or detailed as network provider NPMD tools, it can add decent flavor to the DevSecOps and overall observability.

Elastic Observability is best suited for mid-sized to large enterprises with the resources available to make the most of its features and manage its complexity. Enterprises can deploy Elastic Observability as a SaaS on a public cloud, as a self-managed solution on a public or private cloud, or on-premises. It offers a free tier with full diagnostic functionality, but the useful functions and support require the full version.

Elastic offers a strong open telemetry-based full observability solution that helps enterprises observe the full stack (including APM, infrastructure, services, and the network) of both enterprise and cloud-based solutions. It even provides visibility into combined multi-cloud and hybrid-cloud options within one stack. While the log and metric capabilities are top notch, Elastic should consider expanding their real user monitoring, user experience monitoring, and native mobile application monitoring to make it a complete solution.

Being free and open makes Elastic one of the top choices for cost-conscious customers, and having open telemetry ensures customers can expand and integrate with a suite of solutions coming out of the Cloud Native Computing Foundation.

While Elastic is a strong solution, we would like to see more native integration with top AIOps platforms to combine full stack observability with AIOps. We also would like to see improvements in the alert, incident response, integration with ITSM/DevOps tools, and reporting to make Elastic a solid choice.

Strengths: Having multiple variations to run on-premises and in cloud locations makes Elastic a good choice for enterprises considering multi-cloud or hybrid-cloud deployments. The concept of placing data next to the application and doing federated searches across clusters and clouds is compelling. Low cost of acquisition.

Challenges: Logstash/Beats log collection is effective, but tighter integration with FluentBit for Kubernetes would be welcome. Elastic supports Docker log collection via Beats, but the lack of a native Logstash plug-in creates more work for developers. Maturing the integration with Prometheus would be an improvement for cloud-native solutions, while some customers have reported problems with Kibana-based metrics.

Epsagon

Epsagon is cloud native and uses AI and ML extensively in its solution. Founded in 2017, the relatively new Epsagon service is integrated with AWS and with the AWS Lambda service, making it a good choice for organizations aligned with AWS serverless functions.

Epsagon intends to simplify development and troubleshooting thanks to its lightweight architecture and auto-instrumentation. The platform uses a primary dashboard to give users visibility into complex enterprise architectures. Through that portal, AI and ML algorithms quickly alert users to problems, reducing MTTR.

Using variants of PostgreSQL, MySQL, and Elasticsearch on the back end, Epsagon aggregates, unifies, analyzes, and correlates data from a range of the third-party tools. It aims to deliver a “single pane of glass” for visibility into containers, microservices, Kubernetes, Apache OpenWhisk,serverless architectures, and other infrastructure components. It integrates with ticketing systems such as Jira, GitHub,Clubhouse, ServiceNow, and webhooks.

Epsagon Service Maps automatically visualize all application and infrastructure dependencies to help visualize performance metrics on custom dashboards. Organizations can use these dashboards to monitor an application or to troubleshoot specific areas of that application.

The platform can also visualize application baselines and provide visibility or generate alerts based on metric thresholds at various confidence scoring intervals. It has self-healing capabilities through parsing alert and payload data. Its trace collectors automatically parse and format trace, payload, and metric data. Traces are automatically collected based upon Epsagon agent auto-instrumentation down to the framework level within a service. No need to deploy time-consuming agent deployments.

The platform’s roadmap includes infrastructure events within the context of traces and metrics, custom metrics, RUM, and added language-agent support. This solution is better suited for smaller enterprises looking to increase their observability within any project. Its pricing model makes scaling potentially expensive, but its tools and automations can help smaller operations decrease their MTTR and generate more value from their ongoing projects.

Though a newcomer to the field, Epsagon has built a decent observability solution. Focused primarily on AWS cloud-native applications, it deeply integrates with AWS, including serverless functions that provide full application level tracing, infrastructure monitoring, and metrics. Epsagon needs to improve its capabilities to make it enterprise-production worthy, including integration with enterprise tools, compliance, security, and governance tools. For AWS cloud-native solutions, it is worth a look.

Strengths: Epsagon’s data collectors and instrumentation is lightweight, easy, and quick to deploy. It performs full-fidelity data collection and analysis—not sampling like some other vendors. It also offers end-to-end visibility from applications to microservices, especially for AWS services, and boasts a “time travel” feature to check the prior status of an application/service.

Challenges: Lacks a lot of enterprise functionality, which can make it challenging for larger organizations to adopt. Epsagon could benefit from adding incident resolution, automation/remediation interaction, and features to enable self-healing applications.

IBM

IBM has adopted a cloud strategy that it believes will make it a leader in the hybrid cloud space. By acquiring Red Hat, IBM brought OpenShift into its stable and in turn made its multi-cloud management platform competitive. The observability platform combined with Watson for AIOps is a solid first step. The first iteration of the solution was released in 2019 and IBM is continuing to improve as it goes along. (*Note: We evaluated the solution before the acquisition of Instana by IBM, another vendor in this space. This evaluation does not take into consideration the capabilities of Instana.*)

The platform runs on Red Hat OpenShift and uses an open, hybrid cloud management platform to provide cluster lifecycle management, application lifecycle management, application performance management for traditional and cloud-native applications, event and infrastructure management, and governance and risk management.

Cloud Pak uses industry-leading SQL and noSQL time series databases. The system is set up as a centralized store, but has processing at the edge and on-server. While it's still under development, the roadmap includes AIOps based on IBM's renowned Watson platform.

The solution is based on Thanos and Grafana, which are both widely used. Although IBM claims the solution is built for legacy and cloud-native hybrid modern applications, it performs much stronger in legacy, on-premises environments than cloud-native.

Cloud Pak for Multicloud Management uses webhooks to notify the system of events. The live system uses the IBM Dashboarding solution and PromQL connectors like Grafana for display. The platform includes predictive insights for baselining and anomaly detection, and offers a topology inventory to automatically capture all configuration changes.

The development roadmap includes integrating with Prometheus endpoints and OpenTracing, as well as new AI features via IBM's own Watson AIOps for incident resolution. AI is expected to add new and enhanced capabilities across cost and asset management, governance risk and compliance, and application operations.

Some of the enterprise-specific functions are refreshing to see as many solutions have overlooked those. Policy-based deployment and cluster governance, chargeback metrics to DevOps teams, auto-incident routing to DevSecOps teams, strong observability features for legacy and on-premises applications and middleware are all appealing features for enterprise use.

Their data ingestion capabilities are good, but lag behind those provided by stronger players in the market.

The platform requires Red Hat OpenShift to deploy, which could be an issue for some enterprises, though the license is included. The solution is offered as an on-premises or IBM-managed program.

This could also present issues for enterprises looking to go fully cloud native and have complete control over and co-locate the services in their cloud environment. While data collection for tracing is good and metrics are decent, logging seems to be lagging. It seems weak for cloud-native applications. Ensuring support for OpenTelemetry and OpenTracing in the roadmap could fix some of those issues.

Overall, this is a solid first step in the right direction for Big Blue. The solution is only a year old and its youth shows—it needs to mature in many areas to be worthy of hybrid and multi-cloud enterprise environments. However, IBM has baked enterprise-specific capabilities into the product, so as it evolves it should serve its existing enterprise on-premises client base well.

Strengths: Enterprise focus includes support for legacy middleware, messaging systems, applications, and VMs that will appeal to existing IBM enterprise customers. The Instana acquisition could bring CI/CD pipeline visibility, closed loop DevOps automation, and add Prometheus and OpenTelemetry, though it may take time to get there.

Challenges: IBM Cloud Pak requires a Red Hat OpenShift base to operate, which may be a problem for enterprises that don't want to go that route. The solution lags in cloud-native capabilities and lacks full OpenTelemetry and OpenTracing integrations, though they are on the roadmap.

Logz.io

Logz.io is an Israeli-based company, with a large presence in the U.S. It primarily uses open source technologies and open standards (such as OpenTelemetry) to monitor, log, collect, search, and analyze observability data. In fact, a vast majority of its revenue comes from its observability platform. Logz.io works well with agile, cloud-native customers, most of which are running Kubernetes in production. The company has more than 900 customers, including Siemens, Unity, and ZipRecruiter.

The Logz.io platform has four elements: ELK-based log management, infrastructure monitoring based on Prometheus and Grafana, Jaeger-based distributed tracing, and an ELK-based cloud SIEM. These are fully managed, integrated cloud services for effectively monitoring, troubleshooting, and securing distributed cloud workloads. While the logging solution has been around since 2014, the tracing and infrastructure components are new (released in 2020). The company also just released an open-source synthetic monitoring system using FaaS.

All data ingestion (logs, metrics, and traces) are open standards-based and use open-source tools: FluentD/FluentBit for Kuberbetes, FileBeat for traditional logs, and Jaeger for OpenTracing, and so on. Logz.io is a cloud-based (mostly AWS or Azure, with some running in GCP) multi-tenant solution, which helps the platform perform better at a large scale and lower cost. The recently released mid-tier storage helps customers reduce storage costs by another 30% to 40%. The platform has native integrations with AWS, Azure, and GCP.

The company's focus on open source lets it build functionality on top of established and well-integrated technologies to help customers and provide differentiation within the market. It aims to make the platform easy to use and takes away the need to manage or even understand the setup of the underlying infrastructure.

The platform's AI- and ML-driven Application Insights facility helps roll up, condense, deduplicate, and correlate data. Its Cognitive Insights feature uses pattern matching around errors. It runs queries through search engines to find relevant insights and discussions happening around similar errors and delivers these insights to the customer. This unique solution runs with minimal human interaction after initial training. Plus it uses the powerful Elasticsearch from the ELK stack.

The Logz.io Real Time storage tier provides the platform with highly available, redundant storage, while the "live tailing" option lets users view data as soon it is ingested and follows it through the flow. The platform is based on Kafka for streaming data.

The visible markers on metrics, logs, and tracing let customers see incidents such as deployment, build release, and anything else that might affect production systems. This lets customers see potential changes that could affect the system and trace back to the event.

The Logz.io solution is better suited to smaller- to medium-size technology companies looking for a scalable, flexible solution that allows integration across a wide range of platforms. The solution is available on a "pay as you go," month-to-month model, but customers can arrange access via an enterprise agreement based on infrastructure size. It is available currently only as a cloud-based SaaS solution.

Logz.io is a very strong open-source, OpenTelemetry, cloud-based solution offered as SaaS. While the log management solution and search capabilities are top-notch, distributed tracing and infrastructure monitoring have just been released and need to mature before enterprises can go into full-scale production. Application and services monitoring, metrics for business applications, real-user monitoring, and synthetic monitoring components (an open-source version was just released) are missing. Building such integrations with other solutions can be time consuming.

Strengths: Open-source roots yield lower costs compared to some proprietary solutions. Most cloud developers instrument code with open-source systems, making it easier for ITOps teams to incorporate them into their processes, instead of re-instrumenting services with a different tool set in production.

Challenges: Open-source approach could restrict use of solutions like Fluentbit and FileBeat in some enterprises. The application and services monitoring, real time user monitoring solutions, and digital experience solutions are limited or missing. Distributed tracing is new and needs to mature before being used in mainstream applications.

Micro Focus

Micro Focus is one of the longest-running players in the ITSM/Infrastructure monitoring space. Founded in 1976, the company has a long history of building out its technology portfolio, providing solutions in the DevOps, hybrid IT, security and risk management, and predictive analytics markets. With recent acquisitions of multiple companies (HP Software and Vertica among them) Micro Focus is trying to expand quickly into the IT observability space.

The Operations Bridge product automatically monitors and analyzes the health and performance of multi-cloud and on-premises resources across devices, operating systems, databases, applications, and services on all data types. The platform offers event consolidation and correlation engines, and big data analytics-based noise reduction. It integrates end-to-end service awareness with rule and machine learning-based event correlation capabilities delivered on top of a data lake.

The platform provides monitoring as code, so developers can use APIs to send data to Operations Bridge and configure monitoring along with the code development. It allows for automatic monitoring, analysis, and incident remediation throughout infrastructure and applications. This saves time in monitoring setup and problem remediation. Its AIOps capabilities provide multi-mode correlation for increased first-time root cause identification, improving MTTR.

Operations Bridge uses Micro Focus' own Vertica ML platform, which employs the ITOM Collect Once Store Once (COSO) Vertica data lake to apply ML and advanced analytics on data from multiple sources for a variety of hybrid IT operations use cases. Operations Bridge also provides multiple, customizable dashboard options for different user personas and use cases. The platform simplifies and accelerates automating tasks such as remedial actions, notifications, and recovery procedures.

Most of the recent additions for cloud-native observability capabilities were just released and need to mature with functions and features before they can be seriously considered by digital enterprises.

This solution is best suited to larger organizations and environments with the time, people, and resources to make the most of its ML and AIOps capabilities. The product is available as an on-premises deployment through license purchase or ELA.

While Micro Focus has worked to build on its network, IT infrastructure, and ITSM solutions, the company needs to improve in application (APM), multi-cloud, and cloud-native environments. For legacy users still op-premises, Micro Focus makes a strong case, but there are gaps in logging, application tracing, cloud-native distributed tracing, and CloudOps that the company must address before becoming a strong observability player.

Strengths: The business value dashboard is a cool visualization tool for the non-IT executives to integrate business and social media metrics with application performance. Micro Focus came from automating NOC environments and offers automation/remediation for legacy enterprise IT such as systems and network reboots. It also offers legacy IT runbook automation out of the box.

Challenges: A strong network and systems player, Micro Focus needs to improve significantly in the cloud and application observability areas. The Operations Bridge UI lacks refinement and could use a more modern treatment.

New Relic

New Relic is a San Francisco-based company founded in 2008. Its latest platform, New Relic One, is still undergoing rapid development, with a raft of new features rolled out in November and December 2020. The company primarily targets firms in the technology space, with a focus on forward-thinking organizations looking for innovative solutions to their problems. The company's revenues are increasing, with a 30% increase in 2020 to $600 million.

New Relic One is a cloud-based observability platform that provides application performance management (APM); as well as infrastructure, browser, Real User, synthetics, mobile, and native client monitoring.

The platform provides flexible, dynamic observability of infrastructure environments, from services running in the cloud or on dedicated hosts to containers running in orchestrated environments, including hybrid and multi-cloud setups. With infrastructure monitoring, customers can connect health and performance data of all cloud-based or on-premises hosts to application context, logs, and configuration changes.

Cloud integrations collect data from cloud services and accounts, with no installation process required. Customers simply connect their New Relic account to their cloud provider account.

New Relic One offers integrations with Amazon Web Services, Google Cloud Platform, and Microsoft Azure. New Relic connects these cloud providers to its Telemetry Data Platform, Full-Stack Observability, and Applied Intelligence products.

The platform supports hybrid environments and container orchestration systems, including Kubernetes, AWS ECS, AWS EKS, AWS Fargate, Azure Container Service, Google Kubernetes Engine, Anthos, PCF, PKS, RedHat Openshift, Rancher, and Docker Swarm.

The Kafka-based architecture is designed to scale. It can ingest two billion data points per second and allows for a degree of latency, as data received can be up to 24 hours old. Customers can send data from the New Relic platform to other solutions for long-term storage and data mining purposes via an API.

The New Relic platform is experiencing rapid and ongoing development, with new features and integrations scheduled for release in early 2021. Pricing has recently been simplified, reducing 11 options down to three with a unified payment structure. New Relic now offers a perpetual free tier, a full-stack-monitoring license that provides access to all modules on a per-user pricing model, a telemetry data platform priced on ingest per GB monthly ($0.25 per GB) that scales, and an annual subscription model.

New Relic is a comprehensive platform that offers a good balance of features, usability, and scalability via its Kafka architecture. The platform has broad appeal, which gives it a host of opportunities in the marketplace, but it seems services and finance are particular targets. The compatibility with practically any data source, coupled with clear visualizations, give development and operations teams all the tools they need to manage and fix performance issues.

Strengths: Intuitive visualization systems provide quick and easy insight into systems. Kubernetes monitoring is one of the best in the market. Tags allow the creation of dashboards with little IT intervention and the platform is built to scale.

Challenges: Despite a number of AI features, New Relic doesn't offer true auto-remediation, which could put pressure on administrators at times of crisis.

Splunk

Splunk has been in the IT monitoring business for more than 15 years. In 2019, Splunk acquired SignalFx (founded in 2013) and Omnition (founded in 2018), which enhanced the usability of the Splunk platform and transformed it into a full observability solution.

The Splunk solution combines monitoring, troubleshooting, and incident-response solutions that boost application modernization initiatives. Splunk brings together infrastructure monitoring, application performance monitoring, digital experience monitoring, log investigation, and incident response into a single platform.

Splunk’s observability solution includes: Splunk Infrastructure Monitoring (real-time metrics-based monitoring and troubleshooting), Splunk APM (trace-based application monitoring and troubleshooting) Splunk RUM (end-user monitoring and troubleshooting—in beta), Splunk Log Observer (log analysis for DevOps teams—in beta), Splunk On-Call (intelligent and collaborative incident response), and the Splunk Observability Mobile (alerts and charts for real-time updates).

Through the 2019 acquisitions of Rigor and Plumbr, the company also offers Splunk Synthetic Monitoring and Splunk Web Optimization (via Rigor) and bytecode instrumentation (via Plumbr), which is currently in beta. The Service Bureau component that lets service providers view tokens/quotas for visibility and control over usage and allocation is another helpful feature. It also provides shareable best practices, such as dashboard and alerts.

The “NoSample” architecture uses full-fidelity data at all times to give better quality results, and the Real-Time Streaming feature processes all data as it comes in to improve MTTR. The platform uses AI and ML features to aid and guide error identification and correction.

The Splunk solution has a wide range of integrations to aid root cause analysis and resolution, such as

ServiceNow, IBM Z Decision Support, Jira, SolarWinds, Git, Jenkins, Spinnaker, GitLab, and CircleCI. The system offers a broad selection of visualizations and easy access to logs, metrics, and traces to facilitate diagnosis and remediation.

The platform offers simple ways to scale and scale back operations, and makes threat detection and remediation straightforward. The platform is best used by medium-to-large enterprises.

The Splunk Observability Suite is a SaaS offering with a subscription-based model. There are two pricing models—host-based and usage-based—and two editions, Standard and Enterprise. The Enterprise edition offers enhanced AI/ML-driven directed troubleshooting and centralized management capabilities (such as Service Bureau).Splunk On-Call offers user-based pricing and three editions: Starter, Growth and Enterprise. Customers can pay on a monthly or annual basis.

Splunk has emerged as one of the leaders in the observability space with strategic acquisitions and targeted organic solution development. While the “NoSample” data platform offering is good, it can overwhelm some large digital enterprises if they go too deep and collect all possible data without proper tag indexing or reducing the data size.

Splunk is aggressively working to convert existing customers by offering incentives, and the company says it is committed to integrating with the entire ecosystem. Such an effort would help allay concerns we have heard in the field about Splunk’s commitment to interoperability. Splunk’s OpenTelemetry-based instrumentation option (via Omnition) lets cloud-native shops build quickly with an option to convert to enterprise scale.

Strengths: Splunk ingests full-fidelity data from all sources (logs, metrics, and traces) across the full stack. It also provides massive scalability, sophisticated in-stream analytics, and native OpenTelemetry support.

Challenges: Some customers expressed concerns about the cost and potential difficulty of migrating to the Splunk Observability Suite. Even discounted, the overall cost burden of the solution could be above average to expensive in the long run compared to other vendors reviewed.

StackState

Formed in 2015, StackState is headquartered in Utrecht, The Netherlands, and has an office in Boston. This nimble start-up has built its observability data analytics solution from the ground up, and released the current version (4.2) of its solution in December 2020. The company generally provides an update every quarter. It offers some unique features that have found niche applicability with banking and finance customers, mostly in Europe.

StackState is a topology and relationship-based observability solution that maps business services to its applications, infrastructure dependencies, configuration, and changes. The topology relationships are generally pulled from a CMDB storage, such as BMC Remedy, ServiceNow, and the other IT management tools. It collects data by integrating with other third party monitoring tools, such as Splunk, and can be extended with the platform's own agents.

StackState's cross-domain, time-traveling topology improves existing monitoring and problem resolution systems. It provides a cool dashboard that shows not only the current system state, but also lets you time travel to observe systems' previous state. That feature is based on StackGraph, a proprietary versioned distributed graph database, on top of HBase/Hadoop, combined with support for Virtual Telemetry back ends such as Elasticsearch, Splunk, Prometheus, CloudWatch, and Azure Monitor.

The platform has cloud-native support for Kubernetes, EKS, AKS, ECS and OpenShift, Docker, and Docker Swarm. StackState can automatically integrate topology and telemetry from vSphere-based solutions and also has an open SDK that can support legacy solutions.

The platform has a virtualization layer on top of DataLakes, which automatically binds telemetry references to the topology. It supports the OpenTracing, Datadog, AWS Cloudwatch, and Jaeger trace formats.

StackState has a fully autonomous anomaly detection system that takes baselines for all environments. It is able to identify slippage over time and during stressed conditions. It can also trigger actions set for different situations and automated remediations.

The relationship-based topology map dashboard is a cool feature, showing the health of the systems and identifying services not behaving well by showing them in red. The rollup to an affected business service, based on the model pulled from CMDB, helps visualize any affected business systems, such as online banking.

The configuration changes and change management systems can feed information, which is useful when time traveling the topology map to find a specific incident and its cause. The StackState Automated Root Cause Analysis feature pinpoints root cause. It enables examination of data through different "lenses" to manually conduct root cause analysis and suggest or perform remediation.

This solution is a lightweight tool available for either SaaS or on-premises deployment.

A metric analytics solution, StackState integrates with other monitoring tools for metrics information. To become a complete observability platform, StackState needs to improve on other telemetry and signals. Observability is more than just when, it is also about what and how. Different portions of this solution cover observability (topology-based incident identification), AIOps (anomaly detection), and DevOps (change management, CMDB effects by changes to production systems).

Strengths: StackState's persona-based IT services/business services view is nice, displaying views within a specific context while also allowing time travel to view the topology leading to the current state. CI/CD and CMDB integration is intriguing. The system can identify problems caused by changes instead of symptoms so the service can be quickly isolated—useful for DevOps use cases.

Challenges: Needs to improve on telemetry and signals. The efficiency of the system is heavily dependent on the data quality of underlying systems, CMDB, and data providers (such as Splunk, etc).

Sumo Logic

Sumo Logic is a SaaS-based, cloud-native, multi-tenant observability platform. It was originally built as a log management, big data analytics, and SIEM solution, but now Sumo Logic has added tracing and metrics to revamp the product into a full observability platform.

Sumo Logic's Continuous Intelligence Platform™ ingests and analyzes data from applications, infrastructure, security, and IoT sources. It then develops unified, real-time analytics. The platform employs AI/ML to create a smooth user experience when exploring logs, metrics, and traces.

The platform has out-of-the-box integrations for AWS, Telegraf, Kubernetes, and Prometheus. It can use a combination of different back-end databases depending on the deployment, including DynamoDB and S3. Its diagnostic tools are aided by ML and its UI offers multiple ways to access vital information.

Users can create dashboards for each microservice and instantly view all associated metrics, traces, and logs. The visualizations of these elements, assisted by AI processes, allow quick and easy navigation for engineers to diagnose the causes of errors and failures.

This cloud native, multi-tenant service also provides a visual Root Cause Explorer, a Global Search feature, a time and spatial comparison engine, and the Outlier Detection function to continuously compute multiple baselines.

The Sumo Logic Continuous Intelligence Platform solution is best suited for medium-to-large enterprises looking to scale their observability efforts and provide engineering teams a comprehensive solution for quickly and efficiently resolving service interruptions.

The solution uses a unique payment system based around "Cloud Credits." Users buy a number of credits and can access the entire platform by using credits for different features. Costs are relative to the amount of data ingested, and can easily scale up and down, depending on customer requirements.

The Continuous Intelligence Platform also provides new observability features in its log management platform, with an engineer-friendly UI to encourage deeper problem investigation for quick diagnosis. The lack of auto-remediation could be considered a weakness, but its numerous integrations could mitigate this factor.

Strengths: The visual dashboarding of business KPIs and affected metrics identify impacted critical systems and potential revenue loss. Budgeting helps customers specify exactly how much data they want to process and limit consumption. The new flat-pricing model based on data volume could broaden appeal among cost-conscious enterprises.

Challenges: Not optimized for on-premises, private cloud installations. Azure- and GCP-based enterprises, along with heavy serverless users, might find Sumo Logic challenging for data collection and hosting. The solution lacks user/experience monitoring, synthetic monitoring, and mobile native application monitoring and metrics—critical in a full-spectrum observability solution.

VMWare Tanzu Observability

The VMWare Tanzu solution suite, designed to support cloud, hybrid cloud, and containerized applications, now includes its observability platform—Tanzu Observability. VMware is expanding its support for cloud and Kubernetes and this platform, rebranded from the Wavefront product in March 2020, is designed to help produce, maintain, and scale cloud-native applications.

Tanzu Observability is specifically designed to help enterprises with monitoring, observability, and analytics of cloud-native applications and environments. It uses metrics, traces, histograms, span logs, and events. These are aggregated across distributed applications, application services, container services, and public, private, and hybrid cloud infrastructures to build a real-time picture of an entire ecosystem.

Tanzu Observability delivers instant chart rendering and real-time updating, which enables rapid iterative incident triage. Users can create and customize dashboards from a simple widget-enabled tool bench, and dashboards can be self-service scaled to thousands of users across an organization.

Out of the box monitoring integrations include Dynatrace, Grafana, Graphite, Nagios, New Relic, Prometheus, Sensu, Splunk, vRops, and Zabbix. It provides data correlation across applications, infrastructure, developers, and DevOps tools with packages, dashboards, metrics, and alerts.

The platform offers more than 100 analytical functions to navigate and isolate production issues. For anomaly detection, Tanzu Observability uses AI Genie, an AI/ML-based automatic anomaly detection and forecasting tool. Users can create smart alerts to filter non-critical events and capture anomalies in different environments and time periods. Users can also display and correlate events (deployment activities, starting of the maintenance window, and so on) and alerts as overlays on any metric visualization.

Tanzu Observability is a solution best suited for large organizations, and is used by many services organizations. Pricing is consumption-based, using the monthly rate of metric data delivered. This gives customers the flexibility to start with any application size and scale up or down as needed. It is not dependent on the number of hosts or the number of users.

This is a strong offering from VMWare for a full observability suite. While the solution has strong enterprise features such as governance, security, policy, and compliance, it needs to catch up with some of the new digital observability platforms to offer more for cloud-native solutions. It could also improve on chargebacks and deployment policy controls based on compliance. On a similar note, auto-instrumentation or faster instrumentation would be a welcome addition.

Strengths: VMWare is one of the few vendors that treat VMs, enterprise software, and cloud-native solutions as first-class citizens. It is easier to build a true multi-cloud observability system using VMWare. It also provides observability across any Kubernetes clusters running across any type of cloud AWS, Azure, GCP, and on-premises infrastructures.

Challenges: The proof-of-concept and implementation cycles seem to be longer for many enterprises. Time to value is higher than other vendors in the space. It is also one of the more expensive solutions in this space. Most other vendors do pricing adjustments to stay competitive.

Zebrium

Zebrium, based in Santa Clara, CA, was founded in 2017 and launched its GA services in 2020, making it a fairly new entrant to the observability market. Its primary use case is automated root cause analysis and incident detection using AI/ML on logs and metrics so the IT Ops teams can achieve reduced MTTR.

Zebrium is an AIOps/observability platform that uses unsupervised machine learning to auto-detect software problems and automatically find root cause. The system doesn’t require a manual setup. It trains itself on enterprise topology, baselining the system, and is ready to perform incident detection within a day.

The system not only helps human users diagnose the root cause of issues, but also proactively detects problems and sends root cause reports to engineers outlining the problem, where it occurred, and when. Unsupervised learning also has a feature where a human can accept, mute, or reject the results as true or not. The system then uses reinforcement learning to fine tune its algorithm with no additional manual intervention.

Zebrium works almost exclusively with logs and metrics (no traces supported at this time) to conduct detection and diagnosis. It has a native collector and fully supports Kubernetes (including variants like OpenShift). It also has a native collector for Docker and Linux, and supports log collection via logstash or syslog for Windows, VMware and most other environments. Zebrium has built a Lambda function to forward logs (for something like Amazon Cloudwatch). The platform does not use sampled data and can handle large-scale data acquisition up to the petabyte level.

The platform can ingest and analyze any data, including unstructured data, which is useful for logs in particular. Its auto ML learns the structures of log events and looks for hotspots of abnormally correlated anomalous patterns to detect real incidents over basic anomalous signals. By correlating logs from multiple sources with metric information, Zebrium can proactively create real incident reports, with details of potential root cause, without manual intervention.

Using Zebrium technology, typical customers are able to drive down the MTTR for a software incident from hours to minutes. It uses ELK stack for search, data collection, and log management.

Zebrium is delivered as multi-tenant SaaS or can be deployed on a customer owned VPC or can be installed on-premises. It is free to try, and setup is fairly easy and quick. The solution is priced on data volume and the platform can be installed on-premises for larger customers.

The platform can work with any IT monitoring product currently on the market, plugging into its solution and automatically reporting problems and details of root cause to the NOC or SOC.

This is a great first step from a cool start-up—using AI/ML to differentiate real incidents from the anomaly stack. Currently, it only uses logs and metrics, but is planning to add traces in the future.

Strengths: Zerbrium's ML automatically finds the root cause of incidents without requiring manual hunting through logs and metrics. It can also proactively detect anomalies across logs and metrics without requiring human-built rules. In addition, it can potentially reduce "alert fatigue" by surfacing only the incidents that the system thinks are real.

Challenges: While the system is reported to be highly accurate, there is the potential for false negatives that fail to flag real incidents, and false positives are a risk when a change occurs in the system. For this AI/ML system to work properly, both signals (logs and metrics) need to be strong. Only two use cases are currently supported—automatically finding root cause for incidents and proactively detecting new incidents.

6. Analyst's Take

Assessing an emerging market segment tool like cloud observability is always a daunting task, especially when there is a diversity of solutions addressing cloud-only, hybrid-cloud, and multi-cloud approaches. As we separated the wheat from the chaff, what we found was a robust and fast-evolving market landscape occupied by solutions defined by their strengths and weaknesses. While most vendors offer strong solutions in this space, a lot of these futuristic features are still nascent to evolving.

- **Auto remediation/self-healing:** The ultimate goal of observability is to reduce downtime, build resilient systems, and drive to near-zero MTTR. Identifying and delivering insight is not enough. Solutions must automate at least basic remediation to eliminate costly human intervention.

- **Open connectivity:** We found no one solution can optimally address the entire gamut of observability needs. As such, we find that support for open connectivity—across telemetry collection, tracing, automation, and more—is a vital asset that allows organizations to build best-of-the-breed platforms.

- **DevOps and CI/CD integration:** We look forward to seeing tighter integration with the DevOps cycle and tooling, reflecting the fact that code and configuration changes account for the majority of the incidents in production environments. Risk analysis of code changes before deployment, auto rollback of canary/blue-green environments, and potential outage and lost business opportunity analysis are all good targets for future observability systems.

- **What-if analysis and stress testing:** Another good improvement can be to do a what if analysis and stress testing combined with testing tools before deployment.

- **Cost of ownership:** A key pain point for enterprises adopting observability platforms, cost of ownership is negatively impacted by vendor pricing models that too often are not transparent. Lock-in concerns are also high when committing to a single platform. To resolve this, we hope to see more open integration, open connectivity, and open telemetry, as well as consumption-based pricing potentially replacing pricing based on hosts, agents, and systems.

- **Use of real AI technologies:** In our analysis, most vendors use basic ML models for observability and for AIOps use cases. More sophisticated AI technologies such as Natural Language Processing (NLP), deep learning, neural network capabilities, supervised/semi-supervised learning, and transfer learning all offer promise, though their employment is still nascent. Another target of opportunity is unstructured log analysis, where most observability systems today struggle—a common issue with ability to read unstructured application- and service-level logs.

- **Business use case:** While almost every vendor is working from an IT use case perspective, few concentrate on the business KPIs. A business user should be able to view a business application-level roll up and define criticality, opportunity cost, and allowable costs and constraints.There is a lot of work to be done, but it is very exciting to see many vendors work toward a total observability solution rather than a siloed approach. And we see a great deal of creative energy and innovation coming from newer vendors that are defining the use cases and usability requirements to advance the state of the art.

Ultimately, the cloud creates many opportunities for youthful, nimble, agile, cost-effective vendors to compete at the same level as large, established providers. While some of these new players lack the maturity, security, governance, and enterprise-grade character of established solutions, they make up the difference with focused functionality and inspired approaches. The key, as ever, is to assess and select the solution that best fits your specific needs.

7. About Andy Thurai

Andy Thurai is an accomplished IT executive, strategist, advisor, and evangelist with 25-plus years of experience in executive, technical, and architectural leadership positions at companies such as IBM, Intel, BMC, Nortel, and Oracle. He also advises many start-ups. He has been a keynote speaker at major conferences and served as host for many webcasts, podcasts, webinars, and video chats. Andy has written more than 100 articles on emerging technology topics for publications such as Forbes, The New Stack, AI World, VentureBeat, and Wired magazine.

Andy's topics of interest and expertise include AIOps, ITOps, observability, artificial intelligence, machine learning, cloud, edge, and other enterprise software. His strength is selling technology to the CxO audience with value proposition rather than a technology pitch.

You can find more details and samples of Andy's work on his website at www.thefieldcto.com

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