Ovum
TMT intelligence | informa

Market Radar: Cloud-native Application Performance Management

Ensuring optimal performance of cloud applications and systems

Publication Date: 07 Sep 2017 | Product code: IT0014-003329

Michael Azoff

Summary

Catalyst

Application performance management (APM) is an IT discipline that continues to evolve as new technologies appear. The new wave of microservices architecture and containerization has led to what is referred to as cloud-native development, and this approach has one of the fastest adoption rates of any new technology we have witnessed. With cloud-native environments there is the need to monitor and manage these new types of applications, but the challenge is that they are quite different from traditional monolithic systems. This report delves into what the differences are, how these affect how they are monitored, and how a new generation of vendors have stepped up and introduced novel solutions to meet these challenges. We also examine the incumbent leaders in the APM market that have transformed their solutions to expand and cover cloud-native environments. The technical capabilities of these solutions are compared side by side in the Ovum Rainbow Map (a reverse heat map) to help guide what these solutions offer.

Ovum view

Cloud-native development is driven by multiple factors. DevOps continuous delivery allows for frequent changes, often small updates, to be made to live systems. However, with legacy monolithic applications it is often inconvenient to shut down systems for updates, whereas with loosely coupled microservices, changes can be made in live production. Microservices bring many new benefits to developers, such as, for example, allowing the scaling of development effort without hitting the infamous productivity decline when new team members join. Because it is the technology that disrupters from Silicon Valley adopted or in some cases invented to rapidly grow their business, it makes perfect sense to compete with them by adopting the same technology as part of digital transformation programs.

Cloud-native development uses highly distributed microservices, sometimes carried in thousands of containers, with fast-changing topography, which is a challenge in terms of performance management. The new generation of APM solutions is meeting this challenge with a range of technologies that can reduce the complexity of distributed tracing, time series metrics gathering and analysis, and monitoring container/microservices health and transactions. The use of machine learning to improve automation is another prominent trend. For any organization embarking on building cloud-native systems, these cloud-native APM solutions will be essential for providing the eyes and ears into the highly complex sets of messages that flow between microservices.

Key messages

- Cloud-native will become the dominant development style for greenfield projects in the next two to four years.
- Real-time metrics analysis is a must-have tool for getting directly to problem issues. Distributed tracing is another must-have tool for deeply analyzing latencies between services, and is also useful when metrics do not resolve a problem.

- To manage the sheer scale of components in modern distributed systems, automation is becoming a necessity, and this opens the possibility for machine learning to power this automation.
- APM solutions will need to manage a range of types of applications, from monolithic to microservices, and various shades of hybrid in between.

Recommendations

Recommendations for enterprises

All organizations should consider initiating proof of concepts for cloud-native development if they have not already done so. This style of development is being used not only for new and innovative products and services, but also for modernizing legacy systems. However, cloud-native introduces new complexities and is not an easy style of development. This complexity requires APM to help run these systems. Users from every type of organization have a wide range of solutions to choose from in APM, from open source to premium solutions designed for running enterprise mission-critical systems in production.

Recommendations for vendors

All APM vendors must be aware of the changes in the field that the latest waves of technology are bringing. Modern APM solutions must scale well, and have licensing models that are compatible under these scales, working for both vendor and user. The field of distributed tracing and time series metrics analytics has grown quickly. The next wave of innovation we believe will be in intelligent automation. We hear from nearly all vendors that they are busy doing research in applying AI, with some already in the market with intelligent products. Our advice therefore is to examine how AI can make a difference to your product/service, because your rivals are already doing this.

Defining and exploring cloud-native APM

Definition and characteristics

The new-generation software application architecture based on microservices is changing our perception of previous architectures. Within the microservices view of how objects are coupled, everything that came before microservices is essentially a monolith, even service-oriented architecture. Earlier Ovum reports define microservices and cloud-native development, our new term for the whole set of connected technologies, which we define as the combination of agile development, DevOps delivery, and microservices deployed in containers, all making use of cloud infrastructure. Cloud-native APM is APM targeting these cloud-native applications and systems.

Business value and applications

APM is a well-established and relatively mature IT discipline across legacy and monolithic application architecture. With microservices in production, APM has had to grow and extend to meet the

challenges the new architecture introduces. The key difference is the dynamic nature of the new IT environment with highly distributed systems.

In legacy environments, servers are expected to be long-lived and applications long-running, and only occasionally brought down for updates. This creates a static environment where the connections that exist between components are long-lived, and troubleshooting can rely on agents at fixed points in the environment.

With microservices, the concept of an application is broken into communications and transactions that run between services, where the components (services) and infrastructure (containers, servers, and so on) may appear and disappear solely to support the transitory transactions. Servers are short-lived and are never changed in production. They always exist as source code (infrastructure as code), and changes are made by terminating them in production and rolling out new instances.

In dynamic environments, an issue will therefore be linked to the instance of the environment at the time of the incident, and this continuous change poses a new challenge for APM. The other main challenge in cloud-native environments is the sheer scale of the number of components that need to be monitored. With an application broken into thousands of services, and traffic scaling achieved by horizontally cloning services, the output from metric generators can become a big data problem.

The APM field has evolved to meet these challenges and this report is mainly concerned with these new aspects of the field. Figure 1 shows the range of APM coverage across traditional monolithic applications, to modern microservices, with the central field representing hybrid systems.

Figure 1: Cloud-native APM technology stack

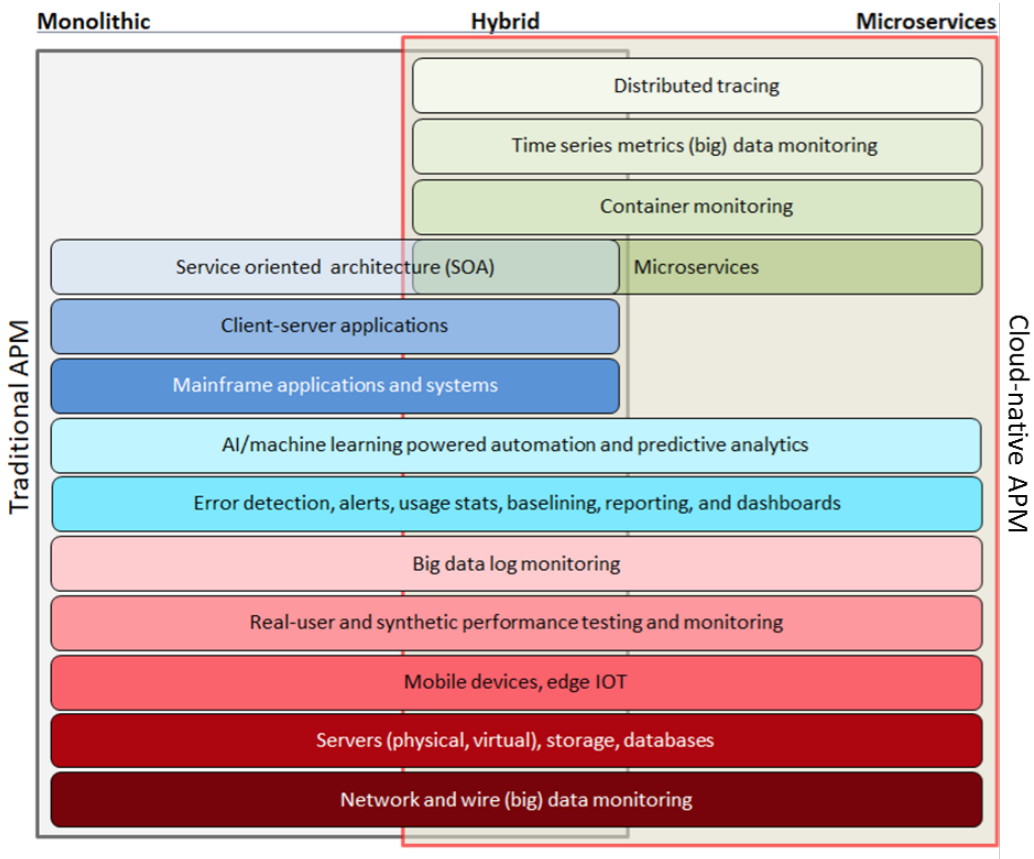

Source: Ovum

Key capabilities

Starting at the top and working down Figure 1, we examine the features that are found in APM solutions. The distinctive cloud-native APM features are those found in the microservices to hybrid column.

Microservices to hybrid systems

Distributed tracing

Microservices architecture and its highly distributed nature pose a particular challenge in performance monitoring, because the scale can become huge and also highly dynamic and transient. A distributed tracing system meets this challenge by monitoring latency issues. It can find and help eliminate abnormal delays to events by gathering the timing data needed to troubleshoot these latency problems. These solutions typically manage both the collection and visualization of latency data stored in a time-series database. Distributed tracing is an essential APM feature for mission-critical microservices systems.

Time series metrics (big) data monitoring

This is a relatively new category that has grown in recent years because cloud-native environments need to handle high volumes of real-time metrics data flowing from agents and taps. Although similar to log data analytics solutions, this category is focused only on metrics, with the capability of managing the data as time series. This category is well served with open source tools, including Grafana, Graphite, StatsD, and Prometheus.

Container monitoring

A good practice is to place one microservice per container (such as Docker) to make deployment very straightforward. Containers can also be used to carry traditional monolithic applications into production, hence the crossover into hybrid use cases. Monitoring containers can be accomplished through the container's own diagnostics, and inspecting traffic going in and out of the container. Solutions that monitor the inside of containers have the potential to go deeper into troubleshooting.

Microservices

Ovum's reports on microservices provide a definition of microservices (see Further reading). Traditional APM solutions struggle to monitor microservices because their volume can swamp metrics gathering, and their topology may be transient.

Monolithic to hybrid systems

Service-oriented architecture (SOA)

In contrast with microservices, SOA offered services that were typically less single-purpose and relied on an enterprise service bus. Previously state-of-the-art SOA is now also classed as monolith by the standards of microservices. Services run a spectrum of degrees of independence, resiliency to down time of other services, and ownership of data, creating a spectrum of macro- to micro- services. Traditional APM solutions serve well here.

Client-server applications

The state of the art a few years ago (now traditional/legacy software applications) were highly coupled components of code, with multi-tiered, client-server architecture. The majority of applications that enterprises will want to monitor will fall into this category. Despite the rapid adoption of microservices,

there are still many use cases where microservices will be considered overkill and a simple client-server application will do the job. There is also the hybrid category of “breaking the monolith”, where legacy systems are modernized by breaking modules for modernizing with microservices. From an APM point of view, the solutions available for this category are mature and plentiful.

Mainframe applications and systems

Mainframes still run the majority of heavy batch processing loads in the financial sector and have not disappeared. However, they are very much part of the legacy heritage of applications and systems, and a number of vendors offer APM solutions that monitor within mainframes not just the input and output. These vendors include CA Technologies and IBM.

Across all types of applications and services

AI/machine learning and predictive analytics

This new category in APM has emerged in recent years. With the resurgence in artificial intelligence (AI) and machine learning (ML), there is added impetus in how vendors are exploiting this technology. This is more than just a nice-to-have technology, such as predictive analytics, and is essential when considering the big data volumes of metric data generated in microservices-type environments. The task of manually making sense of such volumes of data is impossible and automation is essential. Automation powered by ML is being researched by nearly all the vendors engaged in this report, with some vendors, such as Dynatrace and New Relic, already having products in the market. The future of this category is in using ML to move beyond predictive analytics, which aims to pre-empt trouble before users are impacted or notice, and offer automated remedial action (under supervision).

Error detection, alerts, usage stats, baselining, reporting, and dashboards

The workhorse category in APM covers error detection, alerts to operators, system usage statistics gathering, baselining metric data, and reporting and creating dashboards for operators to senior management. The baselining approach of simple straight-line thresholds has now evolved to sophisticated capturing of behavior patterns of applications or systems under normal usage that is used to compare with anomalous behavior.

Log data monitoring

The need to manage big data loads of machine data (also commonly called log data) became recognized when Splunk entered the market, and has since been followed by many other start-ups and incumbent vendors.

Real-user and synthetic performance testing and monitoring

This field has grown with solutions such as crowd-sourcing that can test the applications and devices (mobile apps) across the globe harnessing the internet. Synthetic performance testing providers such as Dynatrace, pingdom, and ThousandEyes can simulate traffic loads sourced anywhere.

Mobile devices and edge Internet of Things (IoT)

Managing performance of mobile apps is a dedicated field available in enterprise mobile management solutions. IoT devices will typically have sensors that are used for diagnosing and troubleshooting issues in host machines.

Servers (physical, virtual), storage, database monitoring

Monitoring servers, storage, and databases is often a specialist activity, and customers will typically have dedicated solutions in place or opt for solution suites that span everything.

Network and wire (big) data monitoring

Application performance is intimately connected with infrastructure performance, and the large IT players with APM solutions will tend to have solution suites that cover the whole spectrum of IT. When looking specifically at application performance, ensure that the APM solution integrates with your infrastructure performance tool (NetScout or Riverbed, for example). Some vendors offer wire data monitoring (data packets flowing through the network) as defined by communications protocols (and therefore distinct from log data) and can handle big data volumes. An example is appliance provider ExtraHop.

Market landscape and participants

Market origin and dynamics

APM solutions in the market continue to evolve, with many open source providers entering the market with solutions to meet the needs of new styles in application development. The current wave of cloud-native development has made the cloud-native APM market highly disruptive. The large enterprise suite providers are re-architecting and expanding their solutions to meet the threat from Silicon Valley. These tend to be highly agile and DevOps in their development cadences, and often use microservices in their architecture.

Key trends in the cloud-native APM market

Ovum expects a significant shakeout across the start-ups racing to establish presence in cloud-native APM. This project approached a select number of vendors of all sizes to participate and is by no means exhaustive. Many players have not been included, such as new start-ups that emerged while the project was in progress. The vendors that have been included are highly representative of what is available in the APM market and include the market leaders.

The biggest trend in the market is supporting cloud-native applications and systems. It is still a nascent sector, and this is the reason we have not produced an Ovum Decision Matrix (ODM) but have instead opted for a Rainbow Map without ranking vendors. The next few years will see this sector stabilize and we will produce an ODM when it does.

Another significant trend we see is in the use of machine learning. Nearly all the vendors we spoke to had some activity in place, from products still in research, to products already in the market. The recent advances in artificial intelligence have given this feature, which was already active back in 2015, a new impetus. The types of machine learning applications we see include:

- Processing big data amounts of streaming metrics using intelligent automation.
- Predictive analytics: this has been a long-term use case but we expect to see availability of improved solutions.
- Advising operators with remedial action to help solve problems.

Distributed tracing was invented at born-on-the-cloud internet companies such as Google and Twitter. It continues to grow as a specialized feature for trouble-shooting microservices messaging latency problems. The growth of time-series metrics analytics solutions has also been driven by solving complex distributed system problems.

Future market development

Ovum expects consolidation in the market in the next few years. There are still start-ups emerging with novel solutions, so this field will continue to be disrupted. For example, Netsil has emerged in the last few months with a novel software-based solution for gaining application insights and system health metrics by analyzing service interactions. As a result, it can produce an auto-discovered visual map of the connectivity in distributed systems that can be saved in a time-series database without needing to add headers to data packets or instrument code.

The future holds much promise for AI applications. APM is rich in data and this is ideal for AI. We expect to see a broad increase in the use of this technology. For example, the use of AI to assist in analyzing both distributed tracing and time-series metrics is highly likely.

Vendor landscape

The key participating vendors

Ovum invited a selection of APM vendors to participate in our technology review. This included market leaders as well as prominent start-ups with solutions targeting cloud-native APM. The participating vendors were:

- CA Technologies
- Correlsense
- Datadog
- Dynatrace
- Instana
- New Relic
- Rollbar
- Sematext
- SignalFx
- Wavefront by VMware
- Weaveworks

Ovum has written On the Radar reports on some of these vendors, providing analysis of their APM products and go-to-market strategies (see Further reading).

The technical capabilities of the solutions were assessed in a similar way to how Ovum creates its ODM technology dimension, using a deep questionnaire, but in this case we produced the Rainbow Map from the results, highlighting capabilities across different solution segments, and comparing vendor solutions side by side.

Ovum Rainbow Map for cloud-native APM 2017–18

The participating vendors were asked to complete a comprehensive questionnaire (see Appendix) for producing the Ovum Rainbow Map (see Figure 2).

Figure 2: Ovum Rainbow Map 2017–18 for the cloud-native APM market

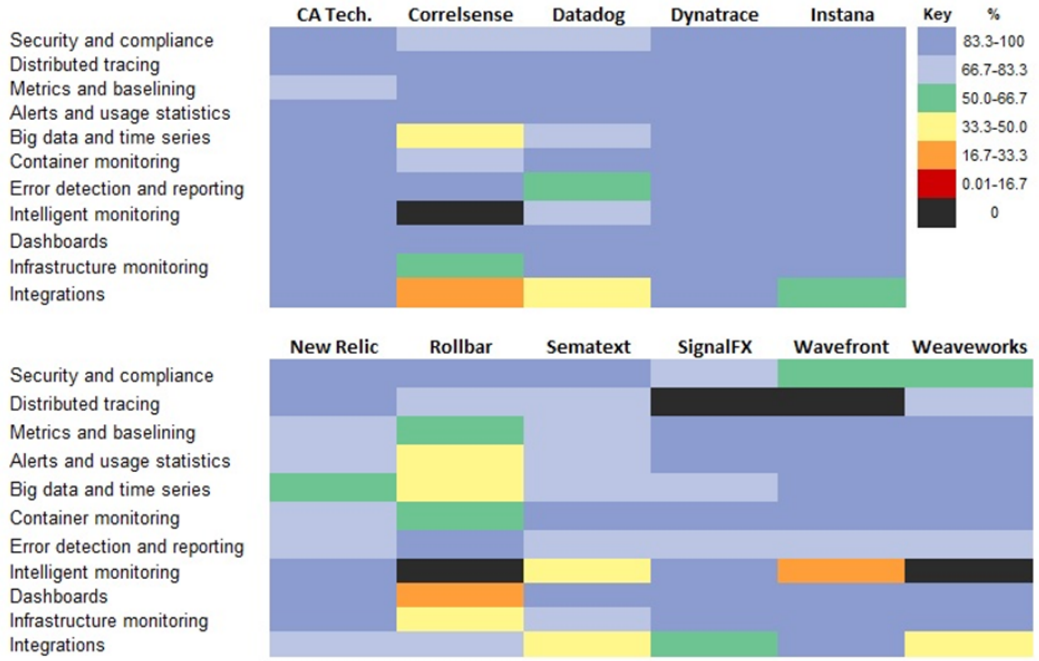

Source: Ovum

The technology sections span APM with emphasis on monitoring cloud-native environments:

- **Security and compliance**: Role-based access and regulations supported.
- **Distributed tracing**: Features to support measuring time lags across transactions. Key feature is viewing all transactions with an error or ID/tag.
- **Metrics and baselining**: Managing metrics, grouping metrics, searching metrics, and baselining using patterns of metrics and infrastructure behavior.
- **Alerts and usage statistics**: How alerts are triggered and types of statistics on system health.
- **Big data and time series**: Analyzing machine-generated data and managing metrics time series.
- **Container monitoring**: Monitoring containers and container orchestration.
- **Error detection and reporting**: Source code level detection to reporting options.
- **Intelligent monitoring**: Use of machine learning for anomaly detection to predictive analytics.
- **Dashboards**: Pre-built and custom options.
- **Infrastructure monitoring**: Any infrastructure performance management available.
- **Integrations**: Integrations with ALM, APM, and collaboration tools.

Other vendors to consider

Vendors not included in this report but nevertheless worth considering in drawing up longlists and shortlists are: AWS, AppDynamics, BMC, IBM, HPE, ExtraHop, Lightstep, Red Hat, Nastel, Netsil, NetScout, Riverbed, Solarwinds, and Sysdig.

Appendix

Methodology

How the Rainbow Map is created for Ovum Market Radar reports

The Rainbow Map is an inverted heat map: the violet/blue end represents the highest scores and the red-end represents low scores. Black is an exact zero score and is interpreted as an area the vendor does not cover. The scores are based on the entries in the Ovum Rainbow Map questionnaire spreadsheet. A number of entries which are for information purposes only are removed, and everything remaining is scored between 0 and 1, and entries within a topic section are weighted. Where a list is requested, e.g. number of tools solution integrates with, the vendor with the longest list is normalized to score 1 and the other vendors are scored relative to that.

Rainbow Map questionnaire with weights

Security and compliance

Role based access to solution 3

Compliant with HIPAA 1

Compliant with HITECH Act 1

List other compliance regulations supported (list) 1

Distributed tracing

View all transactions 1

Can add correlation ID or tag to messages for tracking purposes (e.g. customer ID) 1

Filter view of transactions by: Error 2

ID or tag 2

Map view of communications between services 2

Trace map shows interactions of all elements in infrastructure (servers, databases etc.) 1

Identifies the root cause of an error with pin point of initial service that was faulty 2

Custom tracing support: API 1

SDK 1

Group traces by: Platform 1

Component 1

Service 1

Time period 1

Trace analytics 1

Traffic capture: Sample 1

100% all traffic 2

Metrics and baselining

Create custom metrics (with full monitoring support): API 1

SDK 1

Histogram metrics with ability to aggregate percentiles 1

Support metrics of high velocity (sub-second) events 1

Search metrics data: Query analytics 1

Pre-built search functions 1

Tagging for metrics and other artifacts:

Assigning tags as part of sending metrics 1

Assigning tags in the UI to any metric or object 1

Can send notification messages 1

Can set metrics against individual servers? 1

Can create custom metrics, e.g. revenue or cost 1

Dynamic grouping and filtering (e.g. as maximum and minimum changes the graph adjusts) 2

Group metrics by Zone 1

Region 1

Instance type 1

Image 1

Service 1

Cluster 1

Any custom tag 1

Pattern baselining:

Baselining using patterns of past application or infrastructure behavior 2

State degree of granularity in pattern baselining (shortest time interval) 2

Can pattern baseline be constructed on any analytics query? 1

Alerts and usage statistics

Usage stats: CPU 1

GPU 1

Memory 1

Storage disc activity 1

Garbage collection (GC) Total time spent on GC 1

Frequency of full GC cycles 1

Frequency of incremental GC cycles 1

Event Loop performance 1

Built-in content (dashboards, charts, alerts) 2

Ability to calibrate and preview alert conditions based on real data 2

Detection of alert conditions within seconds 1

Alerts triggered by thresholds set against metrics

Supports dynamic thresholds calculated in real time (versus scheduled or batch calculation) 1

Supports duration requirements 1

Built-in noise dampening 1

Quick problem detection 1

Outlier detection 1

Growth prediction adaptive thresholds 1

Ability to mute alerts 1

Big data and time series support

For how long is the data sampled stored (state maximum time possible for data that is instantly accessible) 2

Analyzing machine-generated big data for: Log data 1

Network traffic data 1

Metrics data 1

Tracing data 1

Other (list) 1

Time series support:

How many open time series can the solution monitor simultaneously? 2

What is the maximum number of metrics per second the solution can monitor? 2

Can perform joins across different time series? 1

Create composite/derived metrics (e.g. ratios) from time series? 1

Can compare values across different time series? 1

Supports ready-built predictive functions (e.g. Holt-Winter)? (list) 1

Can perform standard mathematical functions on data points? (e.g. rolling mean, standard deviation) 1

Perform distribution analytics (i.e. percentiles) on time series? 1

Containers monitoring

Container support: Uses I/O metrics supplied by container 1

Has metrics agent inside container 1

Can monitor container without adding to container image 1

Can pick up new containers - can scale 1

Alert effectively in highly dynamic container environment: How often is alert logic evaluated? 1

Container orchestration management support: Automatic service discovery 1

Maintains context via grouping and filtering 1

Visibility through the stack (for example, linking application to container to host to cluster) 1

Error detection and reporting

Has visualization UI 1

Can identify source code by line responsible for errors or exceptions 2

Data sources used: Performance metrics 1

Stack traces 1

Request parameters 1

URL 1

Environment 1

Users affected 1

Other (list) 1

Report by: Occurrence 1

Browser 1

Deployment 1

OS 1

Location 1

User 1

Host 1

Other (list) 1

Intelligent monitoring using artificial intelligence and machine learning

Detecting unusual metrics behavior / anomaly detection 1

Predictive analytics to identify potential problems and risks 1

Makes recommendations for remedial action on problem detection 1

Dashboards

Secured - require permissions to change 1

Supports custom design 1

Clone/import/export/share dashboards (4 options - indicate which ones) 1

Pre-built content (dashboards, alerts) for integrations 1

Dynamic filtering, time range selection, panning for all dashboards 2

Infrastructure monitoring

Monitors whole infrastructure? 1

List the infrastructure components monitored (e.g. servers, storage, virtual machines) 1

Correlate events across applications and infrastructure 1

Consume infrastructure performance metrics provided by public cloud provider when hosted on its cloud 1

Integrations

APM tools (please list) 1

ALM tools (list) 1

Collaboration tools (list) 1

Notification tools (list) 1

Single sign-on tools (list) 1

Others (list) 1

Further reading

2017 Trends to Watch: Beyond DevOps to Cloud-native, IT0022-000818 (Nov 2016).

Strategy Guide for Enterprises on the Serverless Computing IT Wave, IT0022-000763 (Sep 2016).

Strategy Guide for Enterprises on the Container IT Wave, IT0022-000711 (Aug 2016).

Strategy Guide to the Microservices Architecture IT Wave, IT0022-000702 (Jul 2016).

On the Radar: New Relic Digital Intelligence Platform provides full-stack instrumentation, IT0022-001048 (August 2017).

On the Radar: Rollbar provides error monitoring and analytics, IT0022-001020 (July 2017)

On the Radar: Weaveworks Weave Cloud makes cloud-native development easier, IT0022-000901 (March 2017)

Pending OTRs: CA Technologies, Correlsense, Datadog, Dynatrace, Instana, Sematext, SignalFx, Wavefront by Vmware.

Authors

Michael Azoff, Principal Analyst, Infrastructure Solutions

michael.azoff@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

Ovum

TMT intelligence | informa

CONTACT US

www.ovum.com

analystsupport@ovum.com

INTERNATIONAL OFFICES

Beijing

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

San Francisco

Sao Paulo

Tokyo