Market Impact Report

Juniper Networks' AppFormix: Intent-Driven Cloud-Scale Infrastructure

EXECUTIVE SUMMARY

Today, we live in a cloud-centric world with cloud-native applications and services reaching hundreds of millions of users globally via massive data centers located around the world. Until recently, the cloud has been the domain of a relatively small number of web-scale giants, cloud computing platforms, cloud-native businesses and global software companies. However, enterprises are now migrating IT applications to hybrid clouds and network service providers are reducing costs and increasing service agility by deploying cloud-scale platforms to support Network Functions Virtualization (NFV).

Cloud-scale infrastructure presents significant operational challenges that arise because of the massive scale, software-driven complexity and highly dynamic nature of applications deployed in run-time environments supported by the Docker, Kubernetes and Openstack frameworks, in which workloads and resources fluctuate constantly.

Traditional monitoring solutions rooted in legacy infrastructure are not well suited to the real-time, full stack monitoring requirements of today's cloud-scale infrastructure. Moreover, existing tools are operator-centric and not intended to directly facilitate automated orchestration in response to constantly changing conditions.

Juniper Networks' AppFormix is an innovative, intent-driven cloud-scale infrastructure management system that leverages the inherently distributed nature of cloud-scale computing infrastructure to streamline analytics processing and rapidly detect when service level agreements (SLAs) for applications or workloads are not being met, then automatically notifies the orchestration layer to take remedial action by reallocating resources or redistributing workloads, if necessary.

KEY FEATURES

- Autonomous, intent-driven infrastructure operation for workload and resource optimization
- Smart agents streamline infrastructure monitoring by applying machine learning to metrics local to each node
- Analytics modules monitor SLAs and correlate anomalies and events across the entire infrastructure
- Policy-driven controller assures pre-defined SLAs by automatically driving the orchestration in response to changing conditions
- Compatible with the leading public cloud environments and popular infrastructure frameworks such as Docker, Kubernetes & OpenStack

CLOUD-SCALE OPERATIONAL CHALLENGES

Data center monitoring dates back to the emergence of client-server computing and the distribution of software applications and data across multiple clients and servers, but infrastructure monitoring as we know it today is rooted in the dot-com boom and build-out of large-scale Internet data centers for search engines, e-commerce companies and website hosting. After the dot-com bubble came the Web 2.0 revolution, enabling users to interact and share content, giving rise to web-scale platforms for user-generated content and social media underpinned by Big Data. At the same time, Salesforce.com was proving the viability of delivering software-as-a-service (SaaS) to business users. By the late 2000's, Amazon's Elastic Compute Cloud (EC2) realized the concept of cloud computing. The cloud had gone mainstream and that changed everything.

Cloud-Scale Infrastructure Goes Mainstream

Today, we live in a cloud-centric world. New cloud-native businesses reach customers exclusively over the Internet, serving hundreds of millions of users globally via massive data centers. Consumers have embraced streaming video delivered via the cloud and mainstream enterprise IT is migrating business critical applications to public, private and hybrid cloud environments. All this has been made possible by advances in networking, computing infrastructure and open source frameworks that facilitate rapid deployment and management of cloud applications and services.

Until recently, cloud-scale infrastructure has been the domain of web-scale giants, cloud computing platforms, cloud-native businesses and global software behemoths such as Microsoft and IBM. However, large enterprises are now building private clouds based on leaf-spine switching architectures interconnecting racks of commodity Intel x86 servers to deploy microservices-based applications in container-based run-time environments supported by open source frameworks such as Docker, Kubernetes and Openstack. Network service providers are reducing costs and increasing agility by delivering services based on Network Functions Virtualization (NFV) and are leveraging cloud-scale infrastructure to build distributed clouds for edge computing data centers that will support high bandwidth, low latency services centered on 5G and Internet of Things (IoT) applications.

Cloud-Scale Operational Challenges

Mainstream adoption of cloud-scale infrastructure requires new tools and methodologies so that enterprises and service providers tools can address the operational challenges driven by these factors:

- **Massive scale**. The leading web-scale giants pioneered hyperscale data centers composed of massive arrays of commodity servers based on Intel x86 multicore processors, all fully-interconnected using a Clos network architecture consisting of high performance leaf and spine switches. Servers based on the latest generation of Intel Xeon processors can support CPUs with more than two dozen cores, with a single rack supporting of hundreds of cores. These networks are characterized by east-west traffic flows between servers in different racks exceeding the north-south traffic flowing in and out of the data center.
- **Software-driven infrastructure**. IT application developers are moving to the DevOps model in which small, agile teams rapidly deploy microservices-based applications using continuous integration and delivery methods that place new demands on application infrastructure. Bare metal servers and virtual machine environments are giving way to new software-driven, container-based run-time

environments based on open source frameworks such as Docker, Kubernetes and OpenStack. Continuous monitoring of applications and infrastructure is integral to DevOps, requiring tools for gaining real-time insights into system performance and behavior.

- **Distributed and dynamic environments**. The open source frameworks allow operators to streamline workflows and automate the deployment of applications in hyperscale production environments, but they introduce new complexities in terms of the highly distributed and dynamic nature of microservices application topologies spanning multiple cores, processors, servers and racks. Container-based applications typically have at least 10 times as many microservices per server than virtual machine applications, and these containers will be spun up and shut down constantly according to the needs of the application and workload demand. The net result is high volatility in terms of constantly shifting loads on servers and processor cores.
- **Virtualization**. Network service providers are using NFV to easily deploy and manage a wide array of virtual network functions (VNFs) running on commodity x86 servers, with OpenStack often serving as the enabling software framework. These virtualized environments require new monitoring tools for VNF deployment to manage workload distribution across the data center infrastructure.

This paper examines the gaps in existing tools for cloud-scale infrastructure monitoring and contrasts these with an innovative approach introduced by Juniper Networks' AppFormix for intent-driven, cloud-scale infrastructure management.

TRADITIONAL MONITORING SOLUTIONS ARE NOT CLOUD-SCALE

Initially driven by mainstream adoption of client-server computing in the 1990's, for more than two decades software vendors (see Table 1) have been building tools for monitoring and managing enterprise and service provider data centers. The dot-com boom fueled the commercial adoption of Linux-based open source software and fostered the creation of free and commercially supported instrastructure monitoring tools by communities of open source developers (see Table 2).

Cloud-Scale Infrastructure Requires Real-Time, Full Stack Monitoring

However, despite the breadth, depth and diversity of available options, traditional infrastructure monitoring tools are not sufficient to address the new operational challenges presented by cloud-scale environments, where operators need full stack monitoring tools spanning multiple layers, as shown in Figure 1. Full stack monitoring at cloud-scale requires real-time visibility, machine learning and analytics to rapidly derive actionable insights that drive streamlined operator workflows and automated orchestration for taking remedial action, typically for the following use cases:

- Anomaly detection and root cause analysis
- Performance monitoring to pinpoint bottlenecks and stranded capacity
- Real-time workload balancing and resource optimization

However, traditional monitoring solutions, based on frameworks similar to the one shown in Figure 2, are not well suited for the scale, scope, diversity and highly dynamic nature of cloud-scale infrastructure.

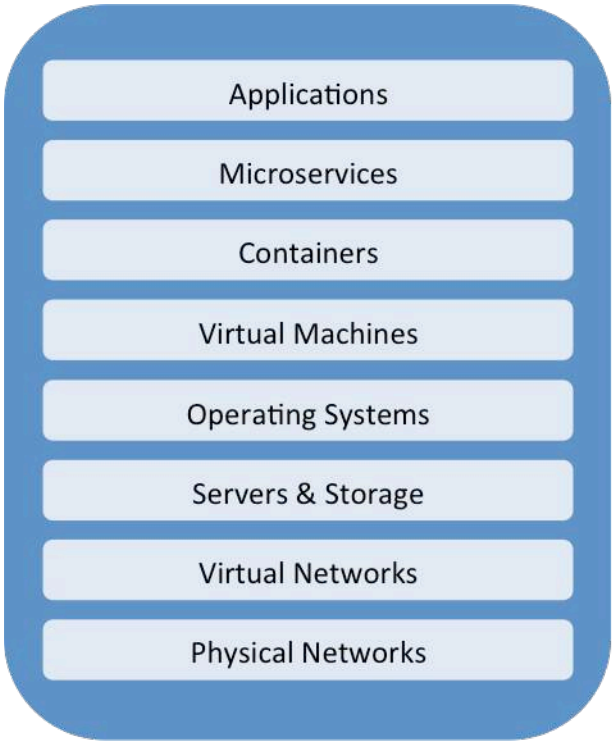

Figure 1. Full Stack Infrastructure Monitoring

To understand why, it is necessary to examine how these systems are designed. At the instrumentation layer, plugins and adapters provide APIs used to extract metrics. Historically, metrics have been pulled using polling methods by a central collector, but more recent implementations are taking advantage of streaming telemetry and message bus mechanisms to push data to a collector. Real-time data analysis typically involves comparing metrics against pre-defined thresholds and generating an alert when a threshold is crossed.

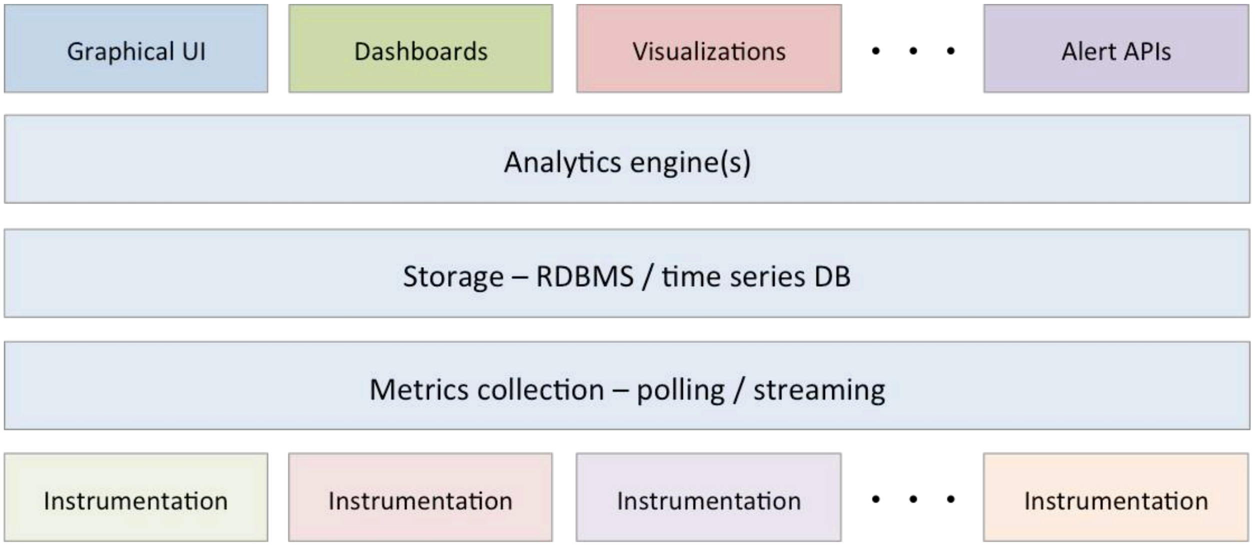

Figure 2. Traditional Monitoring Solution Framework

Retained metrics are stored in a database, historically a RDBMS but more recently a modern time series database. In either case, historical data is analyzed using conventional SQL queries. Operators have access to dashboards and visualizations built using a variety of GUI tools, often vendor-provided, but open source tools are becoming popular for graphical visualizations. APIs for alerts and events facilitate integration with external system management, trouble ticketing and workflow collaboration systems.

Scaling Challenges For Current Monitoring Solutions

Full stack, cloud-scale monitoring requires tracking a wide range of metrics spanning a diverse array of elements and many instances of specific elements (for example, containers and processor cores). Aggregating metrics at a central collection point using polling methods is simply too time consuming to effectively monitor conditions in real-time. On the other hand, streaming metrics to a central collector using a message bus requires a high performance cluster for data ingestion that is usually combined with a Big Data cluster for real-time analytics. Furthermore, alert fatigue is a well-known malady afflicting data center operators, so simply monitoring a large number of metrics to detect when they cross thresholds just increases the noise caused by an incessant stream of alerts and events. Current monitoring methods do not scale gracefully because they either can't track numerous metrics in real-time or if they can, end up increasing the noise level, making it harder for operators to zero in on significant events requiring remedial action.

Current Monitoring Solutions Are Operator-Centric

Existing tools are designed to present information to operators, who must be directly involved in monitoring system behavior and performance, and then accessing different tools and multiple screens to drill down and determine the root cause of a problem when an anomaly is detected. Essentially, operators are mentally correlating event data from multiple sources and then initiating manual workflows in response to changing conditions and anomalies. Scripting tools can help automate these workflows, but operators do the heavy lifting and are an integral part of the feedback loop. It is important to note that this approach also places a premium on operator skills and experience, as well as the number of operators an organization can afford to have on staff. It's one thing for a web-scaler to employ large, highly trained teams of network operators, but this may not be feasible for even large scale enterprise IT organizations or network service providers.

Traditional Monitoring Solutions Are Not Automation-Ready

In large-scale, highly distributed and dynamic workload execution environments, it is problematic to insert human operators into the workflow for anomaly detection and remediation. Under constantly changing conditions and shifting workloads, events can transpire so quickly that relying on existing tools and manually directed workflows can prevent operators from responding quickly enough to address the underlying problem, resulting in outages, performance degradation or inefficient utilization of resources. Ultimately, cloud-scale infrastructure management needs to be insight-driven and automated so that actionable intelligence provides direct feedback for closed-loop automation mechanisms. In addition, automation can also eliminate human errors, which all too often have proven to be the source of major infrastructure outages.

JUNIPER NETWORKS' APPFORMIX

Juniper Networks' AppFormix is a purpose-built, cloud-scale infrastructure monitoring solution that leverages machine learning, analytics and policy-driven intent to enable operators to automate workload and resource orchestration in real-time. The solution takes an innovative, distributed approach to metrics collection that uses local machine learning to dramatically reduce the flow of data that needs to be ingested, analyzed and stored upstream for rapid anomaly detection and root cause analysis. The system rapidly detects when the operational requirements of an application or workload are not being met, according to pre-defined policies, and then automatically triggers remedial action by notifying the orchestration layer. The goal is to take the operator out of the remediation feedback loop and automate infrastructure management based on business intent.

Smart Agent Infrastructure Monitoring

The technical foundation of the Juniper Networks' AppFormix monitoring solution leverages the processing capacity available on each node in highly distributed cloud-scale infrastructure. Nodes based on Intel x86 multicore processors can easily support thin software agents that not only perform periodic measurements to extract a wide range of metrics, these smart agents also locally execute machine learning algorithms on time series data to immediately detect anomalies and outliers, eliminating the need to aggregate a huge volume of metrics at a central collection point for statistical analysis. The net result is that the smart agents extract useful signal data from a noisy stream of nominal operating metrics, detecting not only when thresholds are crossed, but also deviations from dynamic baselines that may not be known up front and may change over time, tracking performance trends and identifying behavioral anomalies. The dramatic reduction in the volume of upstream data reduces compute and storage capacity for data ingestion at the central collection point, and allows for a more lightweight analytics engine and database.

Automatic, Policy-Based Resource Monitoring, Analytics and Control

Figure 3 depicts the Juniper Networks' AppFormix solution architecture. Major components include:

- Smart agents installed on each node that monitor and analyze resource usage
- The DataManager that distributes data from multiple smart agents upstream
- An API-driven, policy-based controller for configuring and controlling the system
- Framework-specific adapters that discover resources and set controller policies
- Analytics modules that correlate and analyze events across the infrastructure
- A web-based dashboard for system and infrastructure visibility

Analytics modules at the core of the system analyze signal data delivered by the data platform, correlating events across the full stack of many elements and layers in the infrastructure. Resources are monitored against policy-based SLAs that are configured during the discovery process. When an analytics module detects a change in conditions that violates an SLA, the controller is triggered to automatically take remedial action, and does this by notifying the framework-specific orchestrator, providing the necessary context so that the orchestrator can automatically adjust workloads or reallocate resources, without operator intervention. The data platform is also compatible with the broad range of Nagios plug-ins for infrastructure metrics. The intent is to allow the platform to be easily installed, configured and integrated into many different operational environments.

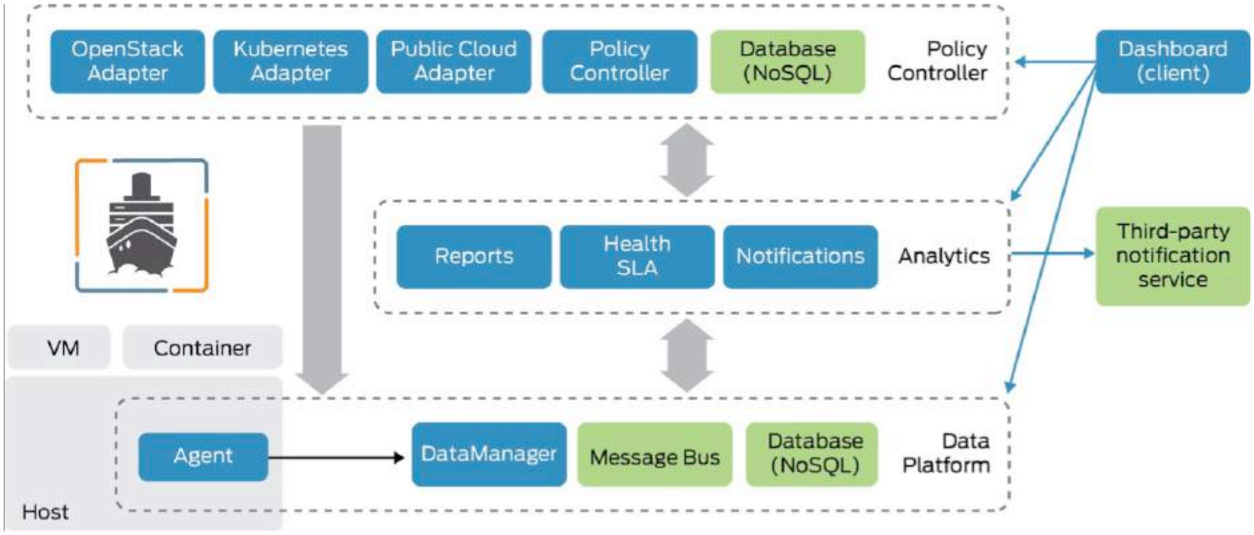

Figure 3. Juniper Networks' AppFormix Solution Architecture

Juniper Networks' AppFormix is designed to run across multiple platforms in public, private and hybrid cloud environments, and currently supports adapters for the Docker, Kubernetes and OpenStack frameworks, as well as adapters for AWS, Azure and VMware environments. A key function of these adapters is to automatically discover and continuously monitor the workload topology, which is constantly changing in dynamic cloud-scale environments. The adapters also incorporate plug-in policies that eliminate the need for operators to configure complex policies for workload monitoring and thresholding.

Juniper Networks' AppFormix enables intent-driven operations for cloud-scale infrastructure, allowing operators to automatically optimize resource utilization in highly complex, dynamic environments, balancing workloads across CPU cores, processors and servers to eliminate bottlenecks and utilize excess capacity. The solution provides the feedback loop between the detection of non-conformance to policy-based SLAs and automatic triggering of remedial actions at the orchestrator layer in accordance with the operator's business needs. The net result is autonomous infrastructure that continuously monitors, heals and optimizes itself in response to changing conditions in order to ensure that customer SLAs are maintained.

Juniper Networks' AppFormix can also export data to other IT operations systems, including central alert and event correlation systems that ingest streams of alerts and events from multiple application, infrastructure and network monitoring tools. Several of these are described below in the competitive landscape section. However, the ability of the solution to significantly reduce alert noise related to cloud-scale metrics monitoring means that these tools can be focused in other areas, such as legacy infrastructure, where the need for event correlation to resolve incidents is still acute.

Business Value of Intent-Driven Infrastructure

Enterprise and service provider data center operators benefit from using Juniper Networks' AppFormix to meet their business and service assurance objectives. Real business value is derived from automated infrastructure management by rapidly responding to incidents, ensuring uptime and reducing the number of highly skilled operators on staff. Automation also helps reduce operator errors and inadvertent outages or non-optimal resource utilization.

Juniper Networks' AppFormix also facilitates capacity planning for optimal infrastructure utilization. The system collects and stores historical data in a modern NoSQL database that can be analyzed to reveal trends and predict future conditions requiring additional resources to assure SLAs. This allows operators to plan in a proactive fashion rather than having to react immediately when capacity is suddenly over-utilized, which often occurs at the most inopportune times.

Historical monitoring metrics are also a source of detailed accounting data for multi-tenant chargeback applications, which is critical for public cloud use cases but also inter-organizational accounting in large enterprises and service providers where data center operations are funded by internal client organizations. Juniper Networks' AppFormix can also be used to provide cloud tenants with visibility into the performance of their applications and workloads, using the same instrumentation and metrics that the operator is leveraging.

COMPETITIVE LANDSCAPE

Juniper Networks' AppFormix has entered a competitive market populated with many established and emerging vendors, as well as numerous open source projects, all focused on integrating machine learning and Big Data analytics into innovative solutions for streamlining data center operations. This section briefly surveys this landscape, starting with traditional monitoring solutions from the leading software vendors and open source projects, as these represent potential competitors, and then highlights new players with competitive products that are incorporating machine learning and Big Data analytics. However, it is significant that across the entire expanse of this market landscape, there appear to be no commercially available solutions comparable to Juniper Networks' AppFormix that are leveraging these technologies combined with policy-based intent to automate cloud-scale infrastructure operations.

Leading Monitoring Vendors

Data center operators have access to a rich set of monitoring solutions from the leading vendors listed in Table 1, encompassing network performance monitoring, infrastructure monitoring and application performance monitoring. However, most of these vendors provide solutions based on products that are 10-20 years old. More recently, VMware has emerged as a provider of monitoring solutions for virtual machine infrastructure, and SolarWinds has leveraged a series of acquisitions to assemble its product portfolio. While all of the companies in this market are applying machine learning and Big Data analytics to multiple areas in IT operations for troubleshooting, application assurance and resource optimization, the solutions remain operator-centric and are not designed for intent-driven, autonomous cloud-scale infrastructure.

Leading Monitoring Vendors
BMC Software
CA Technologies
HP Enterprise
IBM
Microsoft
SevOne
SolarWinds
VMware
Zenoss

Table 1. Leading Vendors with Infrastructure Monitoring Solutions

Open Source Monitoring Solutions

On the open source front, starting 10 to 20 years ago, many developer communities (see the examples listed in Table 2) have created monitoring tools that are freely distributed, with some companies packaging and selling these tools as commercially supported products. Nagios is the oldest and most successful, providing a set of tools and a large number of plug-ins for extracting metrics from many different environments. In contrast, the Prometheus project is a relatively new initiative started only five years ago at SoundCloud, and is focused on a new generation of technologies tailored for monitoring cloud-native infrastructure.

Open Source Monitoring Projects
Calipso
Ceilometer
CollectD
Grafana
Graphite
Heapster
Icinga
InfluxDB
Monasca
Naemon
Nagios
OpenTSDB
Prometheus
Sensu
Shinken
StatsD
Vitrage
Zabbix

Table 2. Open Source Infrastructure Monitoring Projects

For years, data center operators have been successfully tapping into this wealth of open source monitoring software, but in addition to being operator-centric and read only, the design center for these tools is a relatively static run-time environment compared to today's dynamic cloud-scale infrastructure, in which workload topology and utilization is constantly changing. Furthermore, open source monitoring solutions are seldom turnkey and typically require operators to integrate the various components needed into their tool set, so they need to weigh the benefits of customization against the costs of integration and ongoing maintenance, such as incorporating updates and new components as the open source tools evolve.

Big Data Monitoring Framework

Recent advances in Big Data technology, driven by many active open source communities, have enabled the growing adoption of network and application performance monitoring solutions based on the framework shown in Figure 4. Monitoring data is collected from multiple sources, including software instrumentation, network wire data and event logs. Streaming telemetry is the preferred mechanism for aggregating data, which is typically ingested via a high-performance message bus that feeds telemetry data into a Big Data cluster, populated with an array of nodes based on off-the-shelf x86 servers for distributed analysis and storage of data across the cluster. Streaming analytics is usually performed on time series data as it is ingested. Retained data is usually stored in a column-oriented Big Data database well suited to multi-dimensional analytics and capable of responding to complex queries in seconds. Access to data is provided through a rich set of REST APIs, which are used by analytics tools for operator dashboards and visualizations via a graphical user interface.

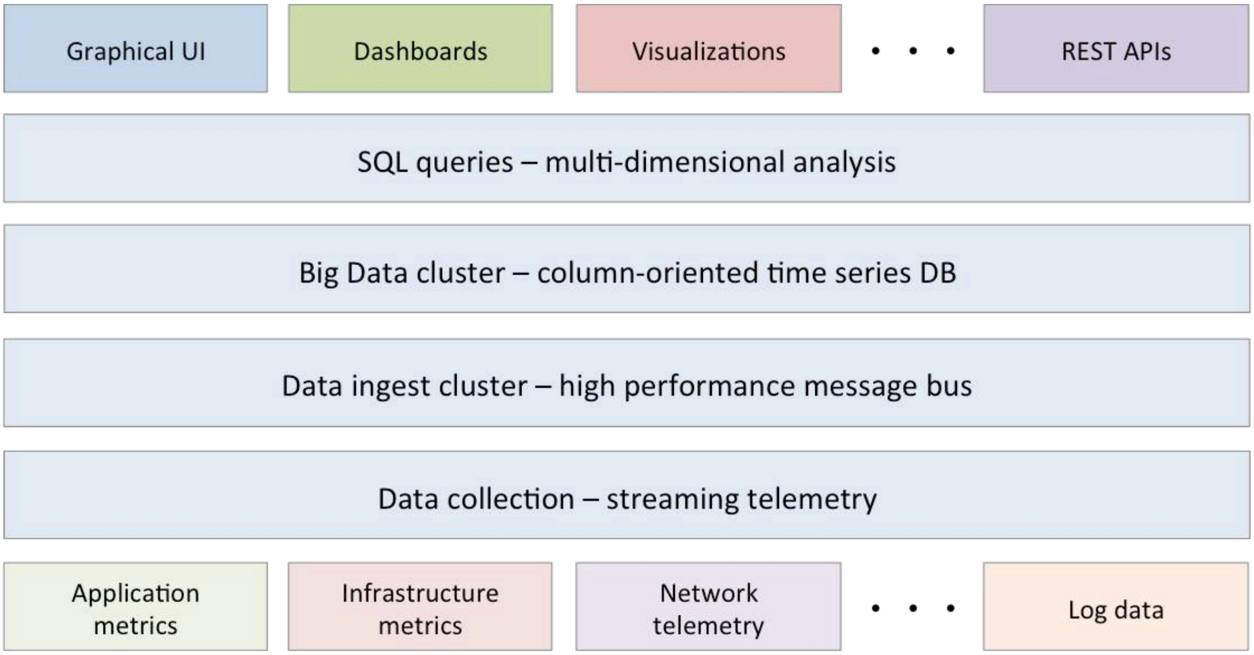
Graphical UI
Dashboards
Visualizations
• • •
REST APIs
SQL queries – multi-dimensional analysis
Big Data cluster – column-oriented time series DB
Data ingest cluster – high performance message bus
Data collection – streaming telemetry
Application metrics
Infrastructure metrics
Network telemetry
• • •
Log data

Figure 4. Big Data Monitoring Framework

The Big Data monitoring approach has been proven effective for ingesting and analyzing huge volumes of network and application performance data in real-time. Big Data analytics can rapidly correlate data from disparate sources to detect performance anomalies and changing conditions at the network and application layers. It is also technically feasible to apply the same approach to collecting and analyzing infrastructure monitoring metrics. However, an approach that aggregates all metrics at a central collection point would not only place significant scaling demands on the message bus and Big Data clusters, it is simply not necessary. Juniper Networks' AppFormix takes a different Big Data approach to infrastructure monitoring that leverages the distributed computing resources readily available in cloud-scale data centers to streamline upstream data processing, allowing for a more lightweight implementation of the analytics engine and time series database.

Open Source Framework Monitoring Projects

The OpenStack and Kubernetes frameworks each take an add-on approach to infrastructure monitoring, and operators are fee to use many different tools. It's notable that Juniper Networks' AppFormix counts managed cloud computing provider and OpenStack pioneer Rackspace as a strategic partner. The Juniper Networks' AppFormix solution is an integrated component to the Rackspace Private Cloud and included as part of its managed service offering.

The OpenStack initiative has five active projects for monitoring and metering:

- Ceilometer – metering and data collection service
- CloudKitty – billing and chargebacks
- Monasca – monitoring
- Aodh – alarming service
- Panko – event, metadata indexing service

Ceilometer agents extract metrics from various elements in the OpenStack environment that are then ingested and analyzed by a central monitoring system – Monasca. This is the traditional infrastructure monitoring approach, although it employs push and pull metrics extraction, a high-speed message bus and a modern Big Data database. The drawback is that all analytics processing is performed at the central aggregation point, requiring significant investment in a cluster that can ingest, analyze and store metrics. Ceilometer and Monasca are also operator-centric and not currently capable of providing feedback directly to the orchestration layer to drive infrastructure automation.

Monitoring in a Kubernetes environment is currently centered on a stack consisting of three add-ons:

- Heapster manages cluster-wide aggregation of monitoring and event data
- InfluxDB serves as the time series database for retained data
- Grafana provides the front-end data visualization and dashboard displays

While this is a modern technology stack, it is still operator-centric with no direct hooks for driving automation. In addition, at this time Heapster only collects metrics pertinent to Kubernetes.

Application, Infrastructure and Network Monitoring Vendors

Industry convergence on cloud-scale infrastructure is driving convergence in vendor solutions for application, infrastructure and network monitoring. Application monitoring vendors are adding infrastructure monitoring (New Relic); infrastructure monitoring vendors are adding application monitoring (Datadog); and network monitoring vendors are adding application monitoring (SevOne). This trend is driven by the critical need for full stack monitoring of cloud-scale infrastructure, as shown in Figure 1.

Datadog is an infrastructure monitoring vendor that has recently extended into application monitoring. Its product is capable of ingesting and analyzing metrics from a wide range of environments. Datadog is using machine learning technology for anomaly detection and outlier detection, and performs advanced statistical analysis of time series data at the central monitoring point. The anomaly detection feature tracks the behavior of metrics that vary over time and determines when metric values fluctuate out of bounds compared to past behavior patterns. Outlier detection is used to determine when a particular element is behaving abnormally compared to similar elements in its group, which is useful for identifying over-utilized or under-utilized resources and reallocating workloads to address the imbalance. However, unlike Juniper Networks' AppFormix, which performs local machine learning on the node from which metrics are being extracted, Datadog requires that all metrics be aggregated at the node where the machine learning algorithms are executed, increasing the processing and storage requirements at the central collection point.

New Relic is a leading application monitoring vendor that has recently expanded into infrastructure monitoring. Its cloud-native monitoring solution ingests metrics in a multi-tenant Big Data database where analytics tools continuously monitor metrics to derive insights that are critical to DevOps teams. New Relic recently introduced a dynamic baseline alerts feature similar to Datadog's anomaly detection feature that monitors metric values against a moving baseline calculated using statistical analysis of historical data. The company also plans to apply machine learning in other areas, including infrastructure monitoring.

SevOne is a leading network monitoring vendor that has added infrastructure monitoring to its product portfolio, driven by the migration of applications to increasingly virtualized environments that place a premium on monitoring the supporting infrastructure in addition to the network itself. SevOne's legacy monitoring solutions are operator-centric and deployed on physical or virtual appliances that can be clustered in order to scale. However, the company recently announced a cloud-native "data platform enabling analytics and automation for cloud and virtual infrastructure management" that includes a cloud-native Big Data analytics engine and customizable tools for visualizations and facilitating operator workflows. This cloud-native platform also includes a high-performance data bus for streaming data to external automation controllers, which is recognition that automation is a key requirement for software-driven, virtualized network and application infrastructure. However, unlike Juniper Networks' AppFormix, the business intent and SLA policies for analytics-driven automation functions would be embedded in these separate controllers and not SevOne's software.

Cisco just acquired Perspica, which will be integrated into its AppDynamics application performance monitoring business. Perspica's mission is to "provide artificial intelligence powered analytics and observability for TechOps and DevOps". Its core technology is a cloud-native Big Data platform capable of ingesting high volumes of application and infrastructure monitoring data in real-time. The platform runs five concurrent machine learning engines for analyzing that data:

- Topology discovery
- Behavioral analysis
- Anomaly detection
- Root cause analysis
- Predictive analytics

Although Perspica initially focused on the alert fatigue problem for speeding root cause analysis, the platform can also ingest a wide range of telemetry, log and topology data. The full stack observability approach and machine learning capabilities are a natural complement to AppDynamics, which is focused primarily on the application layer. However, like New Relic, AppDynamics is mainly focused on application developers and not cloud-scale data center operators.

Machine Learning for Alert Monitoring and Event Correlation

As business operations has become software-driven and enterprise IT operations teams have deployed large-scale data centers to support a myriad of applications, managing the supporting infrastructure has grown increasingly complex. The migration of business-critical applications to cloud-scale infrastructure only raises this complexity. IT operations staff use a diverse set of tools for applications, infrastructure and network monitoring, each with its own user interface. There is no mythical "single pane of glass" that provides a context-rich view of the IT operations environment so that operators can identify patterns, correlate diverse metrics and rapidly determine the severity and root cause of problems.

This is starting to change as new vendors are applying machine learning and Big Data analytics to analyzing and streamlining the incessant flow of alerts from multiple tools, alleviating "alert fatigue" and allowing operators to rapidly detect anomalies, correlate events and perform root cause analysis. Several emerging vendors in this space are briefly described to provide examples of machine learning solutions to one of the fundamental problems that Juniper Networks' AppFormix is helping to alleviate by dramatically reducing the number of monitoring alerts.

BigPanda has developed an "algorithmic event management platform" for unified monitoring visibility, algorithmic alert correlation and smart ticketing. Simply put, the goal is to transform IT alerts into actionable insights. BigPanda studiously avoids using the term machine learning, but claims its algorithmic techniques can achieve 95% reduction in alert noise, allowing operators to focus on "correlated incidents" enriched with data that provides the context needed to take remedial action.

Evanios runs on ServiceNow's cloud-based IT service management platform and solves the same fundamental problem as BigPanda. In addition to applying supervised and unsupervised machine learning techniques to streams of alerts for root cause analysis, Evanios also uses machine learning to predict future events, which are assigned a probability of occurring.

Moogsoft has developed patented machine learning algorithms to reduce the volume of alerts and automate event correlation across multiple tools. Like other vendors in this space, Moogsoft's goal is to reduce the noise that results in alert fatigue by analyzing streams of alerts in real-time to detect problems faster and provide the necessary event context to streamline remedial workflows. Moogsoft's "AIOps" system also supports agile workflows and team rooms that facilitate collaboration for faster incident resolution.

CONCLUSION

Juniper Networks' AppFormix has entered a large and fast growing market with an innovative, unique product offering for intent-driven cloud-scale infrastructure. As the brief survey of the competitive landscape shows, there is no shortage of solutions for infrastructure monitoring and analytics. However, existing tools are very operator-centric and not designed to readily facilitate automation.

Juniper Networks' AppFormix solves the full stack metrics Big Data problem by utilizing distributed computing resources inherent to cloud-scale infrastructure. Machine learning applied locally to metrics extracted from each node dramatically reduces the flow of data upstream to where only signal data is ingested and analyzed. This is a different approach than other infrastructure monitoring tools, which collect and analyze all metrics at a central node.

However, what truly differentiates Juniper Networks' AppFormix is the integration of analytics-driven, policy-based control so that the system automatically drives the orchestration layer to take remedial action and recover from a detected anomaly, or reallocate workloads or resources in response to changing conditions or application demand. The system enables intent-driven cloud-scale infrastructure because it operates by monitoring the infrastructure against policy-based SLAs that are defined to assure the operator's business objectives.

Enterprises and service providers can realize business value from intent-driven operations by using policy-based analytics to drive automated remediation and optimization. They also won't have to maintain a large staff of highly trained operators to manage their cloud-scale infrastructure, which is significant because these personnel are not only highly compensated but often difficult to recruit.

Although Juniper Networks is entering a big arena with AppFormix, the company has the first-mover advantage, which is critical in a market where many players are still primarily focused on operator-centric monitoring solutions incorporating machine learning and Big Data analytics. No doubt, other vendors will eventually follow suit, and it will be interesting to see if new open source monitoring projects emerge that adopt features of the Juniper Networks' AppFormix approach.

Analyst Biography:
Stephen Collins is Principal Analyst at ACG Research, leading the firm's practice in network visibility and analytics. He has more than three decades of networking and telecommunications industry experience across many segments of both the enterprise and service provider markets. Stephen has worked in business and technical organizations for many leading hardware and software infrastructure vendors, serving in executive and managerial roles, including: general manager, VP of marketing, VP of product marketing, VP of business development, product line manager and software engineering manager. He has extensive experience bringing new products to market with technology-driven startups and emerging growth companies as a company founder, member of the senior management team, independent consultant and advisor to early-stage investors.

Stephen is a frequent speaker at industry conferences and has authored numerous articles for trade publications. He holds an M.S. in Computer, Information and Control engineering from the University of Michigan and a B.S. in Computer Systems Engineering, Summa Cum Laude, from the University of Massachusetts, Amherst. He currently serves as an advisor to the ECE department at UMass Dartmouth and also mentors students in technology innovation and entrepreneurship at Brown University.

Authorship: This paper was authored by ACG Research, which is solely responsible for its contents.

Sponsorship: Juniper Networks, November 2017.

ACG Research is an analyst and consulting company that focuses in the networking and telecom space. We offer comprehensive, high-quality, end-to-end business consulting and syndicated research services.