

# Juniper Networks Cloud Analytics and Network Telemetry

The Changing Landscape of Analytics in Virtualized Data Centers with an Application-Centric View

## Challenge

Network administrators are tasked with effectively managing applications running in networks transitioning from physical to virtual, adding significant complexity. Few tools are available that can provide the necessary end-to-end visibility into application performance.

## Solution

Juniper cloud analytics and network telemetry solutions give users an application-centric view of the network, providing real-time monitoring, reporting, and visibility into applications running in physical as well as virtual network environments.

## Benefits

- Application visibility and performance management assist with application and workload placement.
- Capacity planning and optimization detect hotspots and monitor latency and microbursts.
- Troubleshooting and root cause analyses correlate overlay and underlay networks.
- Application bandwidth utilization collects application data traffic for offline analysis.

Today's business environment is driven largely by speed, the dynamic use of IT to gain competitive differentiation and advantage, and time to revenue. The leaders are those businesses that leverage virtualization to get a jump on the competition when rolling out new applications and services.

There are a number of related factors driving virtualization in the modern data center: underutilization of hardware resources and processor capacity resulting in valuable assets sitting unused and gathering dust; operational expenses associated with maintaining these underutilized resources; and virtualization's ability to enable new services in the form of applications. Networks are transitioning from purely physical to a combination of physical and virtual, especially in the data center. Likewise, workloads are rapidly moving from physical to virtual as businesses migrate to the cloud for cost efficiencies.

According to Gartner, organizations had virtualized 70 percent of their server workloads by mid-2014; that number is expected to grow to 86 percent by 2016 as the global cloud market matures into a \$118 billion business this year and \$200 billion by 2018 (IDC). Broadband is also growing, says IDC—doubling every year, in fact, to accommodate worldwide Internet traffic volumes that will triple by the year 2017. With the Internet of Everything, industry experts predict that 50 billion devices will be connecting to the Internet, driving tremendous growth in data networks around the globe.

## The Challenge

Data center consolidation, server virtualization/private cloud, compute-layer virtualization, new application architectures, the shift to dense 10GbE and higher network speeds—all are adding tremendous complexity to today's enterprise network. Virtualization and cloud adoption, in particular, are sweeping the storage and networking worlds, pushing the need for greater business agility and giving rise to new technologies such as application containers.

Clearly, the network is more important than ever when it comes to applications, service delivery, and business operations, presenting IT managers with five key challenges.

**#1 Real-Time Monitoring and Reporting:** The biggest challenge IT managers face today is managing their hybrid physical and virtual environments. In order to achieve operational efficiencies, IT administrators rely heavily on meaningful data collection for the purposes of performance management. This data collection has traditionally been user-driven and collected from individual devices. Limited tools are available that perform real-time monitoring and provide a holistic view of the entire network—physical and virtual.

Without a real-time reporting system, operators can't make informed management decisions. Manual readings taken weeks ago and recorded in a static report are useless in today's complex data centers, where workloads vary widely from day to day or even hour to hour. Many of the available tools still can't provide a cohesive, singular view of applications

and the network; instead, multiple tools are used for each layer—physical and virtual—requiring IT administrators to monitor myriad screens, use multiple commands, and manually correlate all of the information needed to identify, locate, and troubleshoot issues in the physical or virtual network.

Most of the solutions available today:

- Are manually driven, per device
- Provide low-frequency and low-capacity data extraction
- Need to know what you want to know
- Offer limited visibility into tunnels and paths

The lack of comprehensive management tools leads to lower productivity and less business agility—attributes that are required to enable the new services that keep businesses running in such a highly competitive environment.

**#2 Performance Management and Troubleshooting:** Another significant challenge network operators face is gathering relevant performance data for the growing number of applications in use today. Operating and maintaining a physical network is difficult and time-consuming enough; the virtual layers that sit on top of this physical infrastructure make it an even bigger challenge to troubleshoot issues related to application performance. The need to manage, operate, and maintain both physical and virtual layers adds significant complexity to the IT administrator's job and makes it difficult to find and fix application performance problems.

What does this mean to the IT organization?

- Each team uses different tools to troubleshoot the same issue from different angles, providing a fractured and incomplete view of potential problems and resulting in multiple trouble tickets from multiple systems.
- These multiple views into physical and virtual layers each must be “debugged” at the device level and then manually assembled to provide some semblance of a “comprehensive” overview.

More often than not, the network is blamed for performance issues, and network admins are tasked with finding the root cause of the problem. However, it takes time and considerable expertise to track data flows over the various paths through the network to find the trouble source. What happens in such a scenario?

- The lack of visibility and correlation between physical and virtual network layers leads to incomplete diagnostics.
- Statistics collected per network device do not offer a complete picture, nor do they capture ephemeral problems, making it difficult to piece information together to form a comprehensive overview of the problem from an application perspective.
- IT teams—specifically network administrators—are forced to operate in reactive mode.

When there are time-critical problems, workarounds rather than fixes are typically applied to quickly resolve the issue. This is inefficient since it merely masks—rather than fixes—the problem. Meanwhile, finger-pointing continues between the various IT teams.

**#3 Capacity Planning:** While the decision-making challenges associated with capacity planning are unique to each data center, invariably it comes down to aligning the switch capabilities—in terms of CPU cycles and bandwidth utilization on individual devices—to determine how much capacity an organization needs. Since there are no tools that can provide a single, comprehensive view of fault, configuration, accounting, performance, and security management, third-party tools are frequently employed to create a dashboard view of physical and virtual network utilization.

Business success depends on the performance of mission-critical applications and the underlying network infrastructure supporting these applications in terms of bandwidth utilization, delay, jitter, and so on. In order to meet an application's peak bandwidth requirements, businesses will often resort to massive network overprovisioning, resulting in an underutilization of resources during off-peak periods.

**#4 Reducing Operational Expenses:** A continuing challenge that data center managers face today is helping their organizations meet budgetary goals while delivering critical services with fewer personnel and limited resources. By leveraging technologies and processes that increase IT efficiency and maximize existing resources, IT can effectively do more with less—both now and in the future.

**#5 Enabling Analytic Business Applications and Tools:**

Available third-party vendor applications and tools that collect data from different devices throughout a domain have limitations with respect to scale and loss of focus, placing an unwanted burden on the business applications consuming that data.

## The Juniper Networks Cloud Analytics and Telemetry Solutions

Juniper Networks delivers new, innovative solutions and tools designed to help customers overcome these challenges—solutions that provide the flexibility to be fully open and address a variety of business environments. In order to operate, manage, and ensure network availability, it is critical to have visibility into and awareness of what is occurring on the network at any given time. Network analytics and telemetry offer extensive and useful detection capabilities coupled with dedicated analysis systems to collect, trend, and correlate observed network activity.

## Features and Benefits

### Integrated Solution for Improved Monitoring and Reporting

Juniper Networks® Cloud Analytics Engine helps customers move from reactive troubleshooting mode to proactive trouble avoidance mode by collecting real-time application traffic data on the end-to-end switching underlay infrastructure, providing visibility into all devices within a network’s physical and virtual layers. This approach also helps customers, who previously relied on multiple teams to sanitize data to perform root cause analysis, by providing them with a singular, correlated view for better coordination between different teams.

The Cloud Analytics Engine uses network data analysis to improve application performance and availability, collecting real-time data and detecting microbursts typically caused by elephant and mice flows. By detecting microbursts in the network (not just on one device), Cloud Analytics Engine improves application performance and the end-user experience. The solution improves application traffic traceability by simulating the actual application by generating a probe that uses the 5-tuple values of the application. The application probe traverses network devices, collecting and reporting back key performance indicator values such as latency, switching delay, and ingress/egress interface statistics to a central collection server that correlates the data and stores it by time series. The solution also provides open APIs for businesses to develop their own visualization tools. Juniper Networks Junos® Space Network Director can also be used to analyze end-to-end application performance.

### Advanced Application Performance Visibility and Analysis

Junos Space Network Director offers advanced network data analysis features that provide visibility into application flow path analysis and improve the performance and availability of mission-critical applications. Network Director’s advanced virtual machine (VM) analyzer tools provide real-time physical and virtual topology views, track vMotion activity, and maintain complete virtual network inventory. In virtualized overlay networks in VMware Virtual Extensible LAN (VXLAN) environments where switches are unaware of the applications and VM traffic flowing through VXLAN tunnels, these solutions provide full visibility, performance management, and troubleshooting capabilities for physical and virtualized overlay networks in VXLAN environments.

### End-to-End Network Provisioning and Management Applications

Junos Space includes a complete suite of error-free service provisioning tools that provide easy-to-use interfaces for network management and troubleshooting, bridging the gaps between different operational teams.

The Junos Space Network Management platform provides comprehensive management of Juniper devices with broad fault, configuration, accounting, performance, and security management (FCAPS) capabilities.

Cloud Analytics Engine, with its advanced analytics capabilities, helps determine the true load of applications and drives higher network resource utilization.

The table below captures all of the use cases of Cloud Analytics Engine in each area of data center infrastructure.

Table 1: Cloud Analytics use cases

Application	Network Virtualization	Networking Element
<ul style="list-style-type: none"> <li>• App Awareness</li> <li>• App Flow Path Analysis</li> <li>• App Placement</li> <li>• Hadoop/IBM Platform Symphony/VMware Workload Placement</li> <li>• Network Latency                             <ul style="list-style-type: none"> <li>- Per Hop</li> <li>- Switching</li> <li>- End-to-End</li> </ul> </li> <li>• App Network Bandwidth Utilization and SLA Monitoring</li> <li>• App Congestion-based ERSPAN</li> <li>• App Flow Classification (Elephants and Mice)</li> <li>• Policy-based App Flow Monitoring</li> <li>• RESTful API for DevOps Integration</li> </ul>	<ul style="list-style-type: none"> <li>• Overlay and Underlay Network Correlation</li> <li>• Network Functions Virtualization (NFV) Placement (Service Provider Cloud)</li> <li>• VM Trace on Overlay Networks</li> <li>• App Flow Path Analysis on Overlay Network</li> <li>• RESTful API for DevOps Integration</li> </ul>	<ul style="list-style-type: none"> <li>• Granular Data Subscription and Streaming</li> <li>• High-performance API-based HFM ~4K Updates Per Sec</li> <li>• Policy-based UDP or TCP Streaming</li> <li>• Flexible Data Decoration, i.e., JSON, TSV, CVS, GPB, Thrift</li> <li>• Microburst Monitoring</li> <li>• Correlation of App and Congestion</li> </ul>

## Solution Components

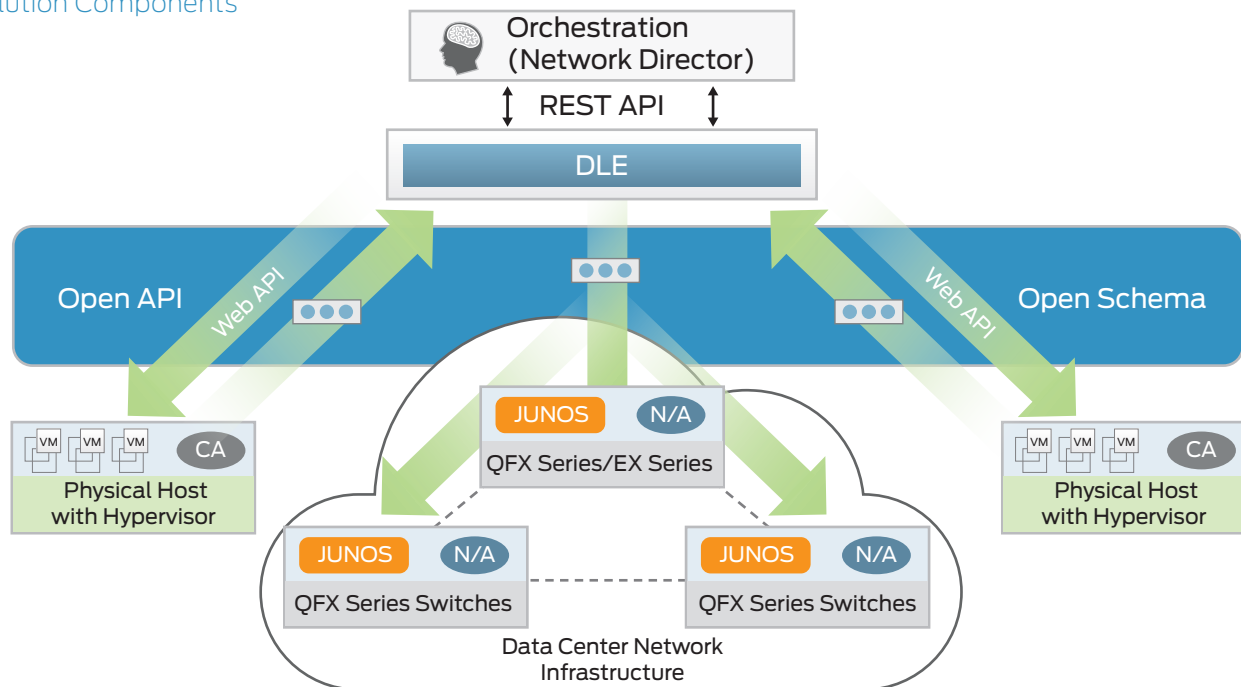


Figure 1: Cloud analytics engine – architecture components

**Compute Agent**, responsible for generating application-specific network data, performs the following primary functions.

- Simulates application/VM or tunnel traffic to collect and store flow statistics:
  - Compute Agent sends probes that simulate application traffic (5-tuple) across the network domain.
  - Network devices that are Cloud Analytics Engine-capable can detect these probes and respond to the Compute Agent with network- and device-specific statistics.
- Identifies active traffic flows on the compute host to create a flow database:
  - Compute Agents detect application flows on the host on which they are running.
  - The orchestration layer or Data Learning Engine can query the Compute Agent to provide a list of active flows on the host.

**Data Learning Engine** adds centralized control and collection for data generated by network devices and Compute Agents. The Data Learning Engine provides RESTful APIs, where the receiver will subscribe to the interest data stream based on supported data entities. The data stream will be UDP-based and will be appropriately tagged with the respective data entities. Application flow streams will be associated with the 5-tuple: a source IP address, destination IP address, protocol, source port, and destination port.

- The Data Learning Engine will support the subscription based on resources such as:
  - Source IP address, destination IP address (for flow data)
  - Network device identifier (for high frequency stats)

The Data Learning Engine will also expose a policy-based TAP interface based on the app-id to provide relevant data for the applications, which can be used by any business intelligence application monitoring to provide the detailed information on application performance and quickly identify issues with application performance issues. The orchestration layer or other external components can also control and access data from the Data Learning Engine using RESTful APIs.

Subscriptions are based on resources. As a result, if the receiver acquires certain additional attributes, those can be discarded at the receiver end; the publisher (the Data Learning Engine) will not filter those attributes. In addition to subscription APIs, the Data Learning Engine will also provide an API to unsubscribe the data. Once subscribed, the unsubscribe API needs to be called in order to stop receiving the data.

**Network Agent (NA)** is included as part of the Juniper Networks Junos operating system package for network switches, which are mainly QFX Series switches for the data center. The Network Agent generates device-specific information like queue latency, high-frequency traffic statistics, and other key diagnostic information. Network Agent does not use any CLI configuration; instead, it provides an API for data streaming. This provides the following benefits:

- Much more scalable than CLI, allowing for rapid configuration changes
- Data streamed in a standard format to an outside collector
- Data that can be generated by Packet Forwarding Engine/ Routing Engine/hardware in any format
- Minimal impact on system

**Junos Space** provides a unified approach to managing a Juniper network infrastructure and for designing and deploying new services. Its centralized management and orchestration for network devices and services is delivered through a single pane of glass for real-time visibility.

**Junos Space Network Director** offers open interfaces and easy programmability. A RESTful implementation of these APIs enables easy consumption and integration with third-party orchestration tools, including CloudStack and OpenStack. It also provides complete end-to-end visibility with application flow path in the data center fabric and bridges the gap between physical and virtual environments.

## Summary—A Complete Set of Cloud Analytics Tools

Juniper cloud analytics and network telemetry solutions provide the complete set of tools required to manage and operate modern data centers efficiently. The Juniper solutions deliver:

- Simple, open and fully integrated APIs at every layer for customers who want to use their own tools to visualize data
- Correlated network and application performance statistics
- Improved IT productivity and user experience
- Visibility into physical and virtual layers, including the overlay tunnels
- Greater efficiency through better capacity planning
- Improved operational efficiency to reduce capital and operational expenses

## Next Steps

To learn more about how Juniper's cloud analytics and network telemetry solutions can benefit your organization, please contact your Juniper representative and visit [www.juniper.net](http://www.juniper.net).

## About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at [www.juniper.net](http://www.juniper.net).

Corporate and Sales Headquarters  
Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, CA 94089 USA  
Phone: 888.JUNIPER (888.586.4737)  
or +1.408.745.2000  
Fax: +1.408.745.2100  
[www.juniper.net](http://www.juniper.net)

APAC and EMEA Headquarters  
Juniper Networks International B.V.  
Boeing Avenue 240  
1119 PZ Schiphol-Rijk  
Amsterdam, The Netherlands  
Phone: +31.0.207.125.700  
Fax: +31.0.207.125.701

Copyright 2015 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos and QFabric are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

**JUNIPER**  
NETWORKS