

Solving the Virtualization Conundrum


Collapsing hierarchical, multi-tiered networks of the past into more compact, resilient, feature rich, two-tiered, leaf-spine or Spline™ networks have clear advantages in the data center. The benefits of more scalable and more stable layer 3 networks far outweigh the challenges this architecture creates. Layer 2 networking fabrics of the past lacked stability and scale. This legacy architecture limited workload size, mobility, and confined virtual workloads to a smaller set of physical servers. As virtualization scaled in the data center, the true limitations of these fabrics quickly surfaced. The economics of workload convergence drives compute density and network scale. Similarly, to meet dynamic needs of business stakeholders, growing data centers must deliver better mobility and its administration must be automated. In essence, virtualized networks must scale, be stable and be programmatically administered!

What the Internet has taught us is that TCP/IP architectures are reliable, fault tolerant, and built to last. Why create yet another fabric when we can leverage open standards and all the benefits that layer 3 networks provide. With this settled, we can work to develop an overlay technology to span layer two networks over a stable IP infrastructure. This is how Virtual eXtensible LAN (VXLAN) was born.

At Arista we work to bring VXLAN to the mainstream by co-authoring the standard with industry virtualization leaders. We're also innovating programmatic services and APIs that automate virtualized workflow management, monitoring, visualization and troubleshooting. VXLAN is designed from the ground up to leverage layer 3 IP underlays and the scale and stability it provides. Let's explore how Arista and VXLAN are solving the Virtualization Conundrum.

EOS®

Just as good architecture is fundamental to stability and utility of a building, the underlying architecture of Arista's EOS (Extensible Operating System) is founded on the requirement to reliably control high performance data center switches in the modern virtualized data center. EOS controls networking platforms that scale data centers supporting thousands of 10G plus ports in a high performance, fault tolerant network. Taking lessons from the Linux world EOS is founded on an unmodified Linux kernel and users gain the stability, security and supportability derived through the open community. From its Linux foundation, EOS delivers users not only these benefits but also extensibility and flexibility.



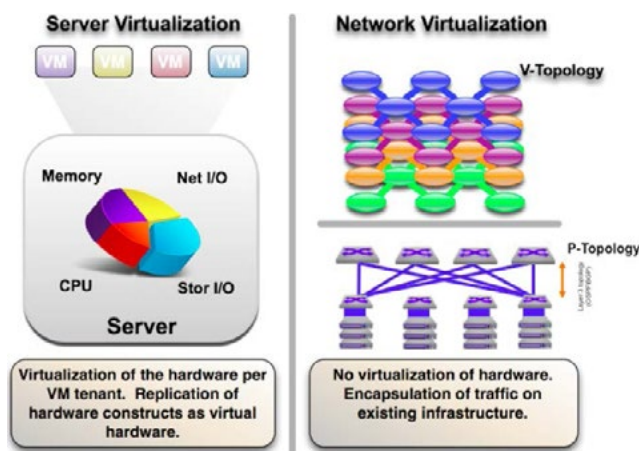
Extensibility shows in every aspect of our innovation in network virtualization. An extension of EOS, Arista's EOS CloudVision® is a platform for workflow automation and workload orchestration. In the context of network virtualization, CloudVision optimizes the virtualization overlay controllers with the physical infrastructure in the modern data center. CloudVision presents OVSDB and JSON-based APIs to leading virtualization controller platforms. It also presents valuable telemetry to monitoring and visualization tools to help improve network visibility and troubleshooting. There's even simple to use programmatic APIs that let administrators or integrators quickly craft rich extensions that improve the manageability and reliability of the virtualized data center. From our suite of troubleshooting tools and virtualization integration points through CloudVision, to our comprehensive APIs for custom extensibility, Arista remains committed to building value, scale and reliability in an open operating system. Arista delivers our customers the most powerful and flexible network virtualization solution.

What is Network Virtualization?

Network virtualization is sometimes confused with server virtualization. However, the two virtualization technologies are quite different yet complementary.

Server virtualization is the partitioning of physical server resources, such as memory, I/O, storage, and CPU. These resources are confined to the physical construct of a single device and share little, if any, distributed state. Network virtualization is the ability to separate, abstract and decouple the physical topology from a 'logical' or 'virtual' topology by using encapsulated tunneling. This logical network topology is often referred to as an 'Overlay Network'.

Overlay networks are necessary to construct containerized workloads whose workflow applications, Virtual Machines and physical hosts are logically interconnected, maintaining the same network IP addresses and subnet characteristics of a 'wire of Ethernet'. The key benefit is workloads can be physically distributed across any resource in the data center thus simplifying scaling, growth and price performance without compromising uptime.



VXLAN (Virtual eXtensible LAN) enables overlay networks on traditional IP topologies. VXLAN is an open, multi-vendor standard that was developed and is being adopted by industry leading network, firewall, WAN optimization and Application Delivery Controller (ADC) vendors. VXLAN's layer 2 tunneling capabilities overcome IP mobility limitations by allowing administrators to extend virtual networks to any VM running on any hypervisor in the data center. This allows administrators to implement a reliable, vendor agnostic L3 architecture in the data center while also supporting VM mobility across all the servers in the data center.

Arista's Network Virtualization

Arista's goal is to provide the best solutions for best in breed overlay networks; to decouple the physical topology from the logical topology – to allow seamless connectivity between compute and network services regardless whether it's physically or virtually connected, and regardless of whether it's local to, or between data centers.

This best of breed approach delivers an optimal level of flexibility, mobility, automation and performance. This ensures workloads can be dynamically placed anywhere, removing traditional layer 3 boundaries of the physical infrastructure.

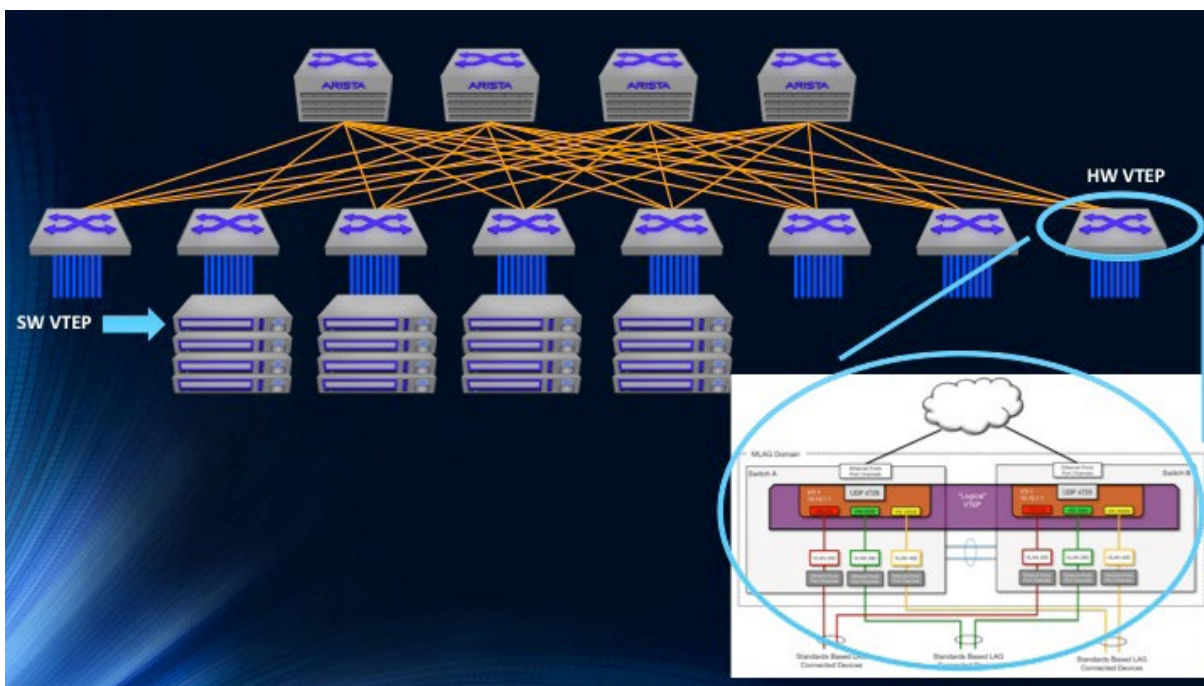
Any vendor can implement VXLAN because it's an open standard. However, to implement a solution that supports business critical workloads, the vendor must provide leading edge virtualization and orchestration services, troubleshooting, visualization and support tools, plus easy to use APIs that support flexible customization and integration of existing customer technologies. Arista EOS and CloudVision were designed with those requirements in mind.

VXLAN Implementation - Integrating Software and Hardware Virtual Tunnel End Points (VTEP)

Arista Networks is proud to be co-authors of the VXLAN specification. Our ongoing contributions to VXLAN technology and to industry leading commercial and open source orchestration offerings have helped this technology evolve into a true solution. Arista's position of being a leader in Software Defined Cloud Networking (SDCN) and the principal authors of VXLAN gives us the unique position of understanding the details of virtualization services and the ability to deliver high performance, production ready virtualization solutions. One of those enhancements we provide is the implementation of hardware encapsulation gateways or VTEPs in our switches. In addition, there's services for emulating broadcast, unknown unicast and multicast (BUM) behavior on the virtual Ethernet segment. BUM services is an essential function for supporting overlay networks but in a way that ensures stability and scale needed for modern day cloud data centers.

VTEPs act as the entry point for connecting hosts, be they physical or virtual devices, into the VXLAN overlay network. The task of the VTEP is to encapsulate local traffic destined for nodes learned on a remote VTEP. In addition to encapsulating overlay traffic, VTEPs play a key role in processing broadcast and multicast traffic and advertising locally learned MAC addresses to neighbor VTEPs. Lastly VTEPs also de-encapsulate traffic from the overlay network and forward it to the relevant locally attached node using standard L2 forwarding mechanisms.

VTEPs are available either as a software virtual switch that is part of the hypervisor software, or as features on physical switches that are part of the underlay-networking infrastructure. Arista ToR (Top of Rack) switches provides line-rate VTEP functionality completely in hardware. L2 encapsulation bridging and routing are implemented in hardware to ensure the best price performance available in a VTEP enabled ToR switch. Building the mapping of device MAC addresses to VTEPs can be either be statically defined, learned via VXLAN multicast services or can be implemented with control plane mechanisms using APIs such as BGP Ethernet Virtual Private Network (EVPN) extensions or mechanisms that are part of the Open Virtual Switch DataBase protocol (OVSDb).



Hardware VTEP as a ToR delivering wire-speed VXLAN to hypervisors and bare metal servers

OVSDb is an industry standard protocol for VTEP management. Among its specifications are communications mechanisms (read / write functionality) for VTEP (software or hardware) to centralized controller / orchestration. Leading virtualization architectures use OVSDb in their integrated virtualized compute and network provisioning solutions. Arista has developed the CloudVision suite of network virtualization services that supports OVSDb and other RESTful northbound APIs.

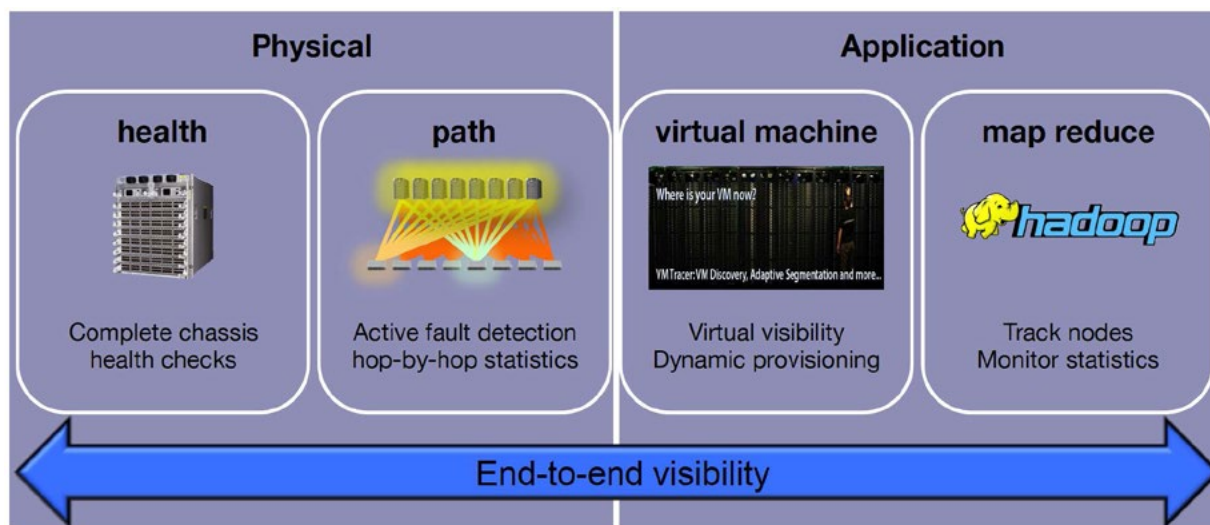
CloudVision is a controller-agnostic platform that can support various virtualization controllers, so the converged data center can support a multi-vendor virtualization architecture on a common network infrastructure, making it more cost effective. CloudVision helps reduce licensing costs and simplifies deployment validation by supporting a single touch-point for the virtualization controller. Controller certification is coupled to the CloudVision release level, not the software running on the switches. This allows administrators to patch infrastructure software without jeopardizing controller functionality. CloudVision uses Arista SysDb southbound APIs that allow scaling to more switches with fewer controllers. Similarly, CloudVision extends the scalability of OpenStack's by integrating with Neutron ML2 plugin services to provide a single touch point for all Arista switches in the data center.

Network Telemetry - Simplifying Operations in a Virtual Overlay Network

As Enterprises and Service Providers evolve from traditional application specific networks to virtualized, on-demand cloud networks, troubleshooting and monitoring toolsets also need to evolve to provide both fine-grained visibility into application performance, and network-wide monitoring capabilities that integrate with both industry standard and customer specific DevOps solutions.

The challenge is to troubleshoot and have visibility when traffic loss happens in an ECMP (equal cost multi-pathing) network. There is an ever-growing need to have end-to-end visibility and monitoring that helps with troubleshooting these large-scale cloud networks from Virtual to Physical infrastructure, and with mission critical applications running on top of the network infrastructure.

As highlighted in the figure below, Arista Tracers are enhanced network telemetry applications that bring deeper visibility by integrating with distributed applications like Big Data, Cloud, and Virtualized environments.



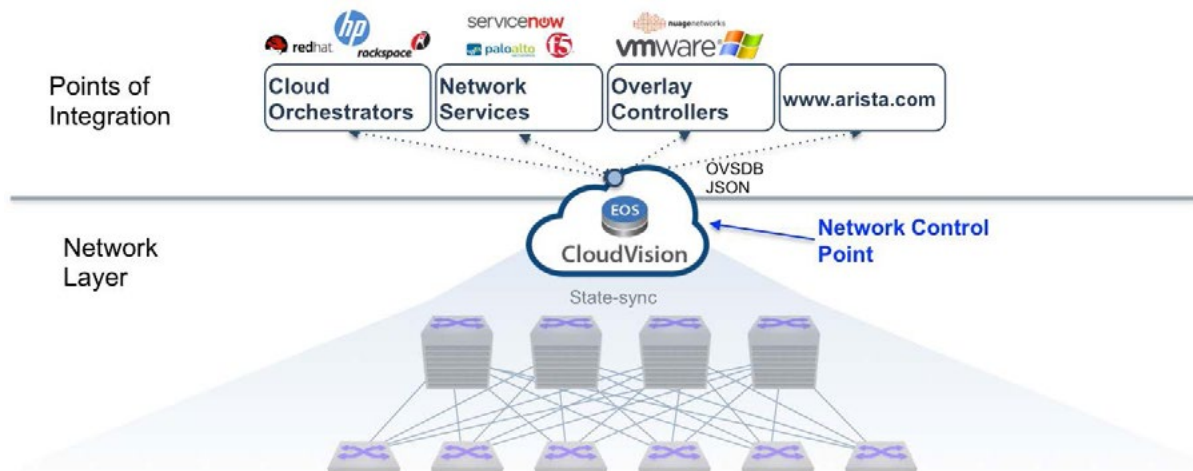
Similarly, in order to enhance visibility into the combined physical+virtual environment, Arista has expanded the scope of CloudVision; providing services that enable full visibility and automated provisioning of virtual environments. These expanded services include:

- VXLAN- Aware VM Tracer - Troubleshooting and identifying a mobile virtual machine is challenging for even a seasoned network operator. If one adds the additional complexity of a L2-over-L3 overlay network architecture with VXLAN, the problem may seem insurmountable.
- VM Tracer revolutionized how network operators identify and track workloads. Within Phase 1, Arista's VM Tracer provides tools to track virtual workloads and their relation to the physical infrastructure. VM Tracer supports overlay visibility, dissolving boundaries between 'overlay and underlay,' so network operators can rapidly identify a virtual machine within a workload, orchestrate it, and auto-provision its network with CloudVision services. VM Tracer supports VMware vSphere, and is being expanded to include VMware NSX-V.
- Arista Path Tracer - Arista Path Tracer probes the physical infrastructure to detect potential intermittent loss conditions that may not immediately degrade application performance or be reported with standard SNMP based tools. These "grey" problems are often elusive in scaled L3 ECMP topologies where many multipath links exist. Path tracer helps identify marginal network conditions more quickly resolve network related performance problems.\

Arista Advanced Event Manager (AEM) technology tracks device and link state changes in a standards based database with open APIs. This allows operators to fully utilize the network state within EOS to automate routine tasks, respond to real-time network or critical device events and take local automated action based to improve network performance and reduce workload downtime. Arista Tracers, coupled with AEM lets administrators script automated tasks utilizing data from rapid telemetry. Instead of using manual, costly and time consuming ticket systems, Administrators can leverage automated scripts to perform routine steps in a virtualized environments. These scripts can be part of virtualization suites, be locally developed, using Arista's easy to use, RESTful JSON API, or can be developed in partnership with Arista's EOS+ Consulting Services (EOS+ CS). Users, for example, can use VM spin-up/down/v-motion to trigger actions without intervention. Arista's CloudVision, AEM and Tracers features are purpose-built to improve system up-time and enable distributed systems automation to simplify complex tasks, react programmatically to network changes, and simplify complex troubleshooting tasks—thus minimizing systems outages.

Partner Ecosystem

Best of breed solutions imply leveraging technology from the industry's most innovative companies. Therefore, building best of breed-virtualized networks cannot be achieved by a single vendor. The best virtualization implementation isn't solely based on multiple vendors using a common service like VXLAN. It also requires collaboration on orchestration and instrumentation services, using open APIs to achieve seamless orchestration, provisioning and monitoring without any vendor lock-in. These properties are at the heart of every successful network virtualization implementation. Arista Networks is the preferred network partners of leading virtualization and orchestration technologies such as VMware, NSX, Nuage, OpenStack, Ansible, Chef, Puppet and many more. Our partner system is among the strongest in the industry and continues to grow because of the power and flexibility of EOS.



CloudVision JSON based APIs deliver a multitude of access points for orchestration and administration

Combining best of breed hardware and software VTEPs allows a virtual network to communicate to the rest of the physical network seamlessly. Today's highly virtualized datacenters require operators to have visibility across both the physical and virtual environments. Many solutions only provide a singular view, which leaves operators blind, limiting mobility as well as troubleshooting. The combined Arista and VMware solution provides true visibility across both physical and virtual. The reason this is possible is because of Arista and VMware's support of the OVSDDB protocol allowing VMware's NSX controller to communicate with an Arista switch for configuration and monitoring purposes.

Arista is also a supporter of OpenStack. Arista is one of the main code contributions to OpenStack's ML2 plugin. As specified by the standard, Open vSwitch has a database schema to save various configuration details of the switch e.g. L2 information, VLANs, port aggregation, QoS, etc. Arista can then also speak OVSDDB with the Open vSwitch residing on the switch (either physical or software) and a centralized controller / orchestration tool. Along with CloudVision and the OpenStack Neutron ML2 and L3 plugins, Arista's networking support for the open standard is arguably the most complete in the industry.

Orchestration and Controller Integration

The role of the network controller has evolved rapidly with the most common implementations moving to a pragmatic model where the network still controls topology construction and the controller provides a single point of integration with northbound orchestration and cloud/virtualization platforms. Actions in the orchestration platform are identified by the controller and notify the network so provisioning can be automated, workloads moved, etc.

Arista EOS CloudVision is designed to be controller agnostic and work with the broadest array of controller vendors. Arista EOS CloudVision integration provides additional topological awareness in the virtualized network. This open and controller agnostic architecture provides the most robust and efficient capabilities in the industry to manage the interaction between the virtual and physical network.

Network virtualization scales better when the controller can leverage network data and control APIs more efficiently. Platforms, such as VMware, Nuage, or open source platforms such as OpenStack actually scale better and run more smoothly when integrated with Arista EOS CloudVision. Reducing controller touch points improves performance and scale by distributing the control load to CloudVision services. Additionally, leveraging these controllers within an Arista EOS CloudVision infrastructure gives customers the advantage of seamless provisioning, segmentation, and policy management of applications anywhere in the network.

Summary

The migration from application specific data center networks to a fully virtualized data center is a multi-faceted discipline that must incorporate multiple virtualized compute platforms, bare metal appliances, physical to virtual monitoring, troubleshooting and automation in the forms of workload provisioning and change management.

Arista's EOS and CloudVision software suite provides a framework of services and APIs that streamline the design of the virtualized data center while also improving its scale, performance and reliability. CloudVision's ability to integrate leading compute virtualization solutions with price/performance leading data center platforms and open extensible APIs helps data center designers implement a cost effective, best of breed data center architecture that can accommodate the requirements of nearly any conceivable business workload. Arista is bringing key virtualization capabilities to the network with a pragmatic and realistic roadmap catered to deliver network virtualization at a useable scale for major enterprises and cloud providers alike.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office

1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. 06/15