Introduction

The rapid adoption of virtualization technologies are driving server consolidation, data center optimization and application mobility. IT organizations are adopting new data center architectures, either in house, or with managed cloud Service Providers. These virtualized data center architectures offer unprecedented asset utilization, with lights out provisioning and automation. Both Enterprise and Service Provider IT organizations need to deploy new network designs capable of massive scale that can accommodate these highly dynamic virtualized workloads and resultant network traffic demands while minimizing complexity and cost.

For many, Software Defined Networking (SDN) is being viewed as the technology advancement to address this new paradigm in data center networking design. SDN promises a shift from distributed intelligence, to centralized intelligence, and from proprietary based architectures to Open Source, including standards based switching and routing solutions. The value propositions of SDN run the gamut including reduced cost of ownership, greater configuration automation, more intelligent end user services, to lower operation costs.

SDN is still an emerging collection of technologies and as the standards that SDN encompasses reach maturity, and hype meets reality, the implementations will be different from the initial concepts. Already, there are several open source software stacks including OpenStack, OpenFlow, and open virtual switch (Open vSwitch/OVS). And there are different tunneling, flow and controller technologies that are coming to market. Further, several companies, are discussing the migration of traditional network management tools to centralized controller architectures. For many network designers, managers, and administrators the variety of SDN architectures can be overwhelming.

Ultimately SDN deployments will be evolutionary, as customers cannot afford to rip out existing infrastructures. In many cases multiple SDN architectures may be deployed within an IT organization to deliver the requirements for scale and operational efficiencies. In other cases SDN architectures may be deployed for failover, application mobility, and capacity management. Selecting a network architecture that can support multiple SDN models ensures IT professionals the flexibility to meet today and tomorrow's business requirements whether they are for Enterprise or Service provide deployments. SDN can be achieved via standard APIs that several hardware and software architectures intend to support.

Arista's Software Defined Cloud Networking (SDCN) Pillars:

Arista Networks fundamentally believes in many of the core SDN concepts- scalability, operational efficiency, and application mobility. Within the Data Center the most pervasive use case for SDN technologies are cloud related, and as a result Arista refers to SDN as Software Defined Cloud Networking (SDCN). The addition of the word cloud by Arista places the SDN focus on networking virtualization, application mobility, and automated service enablement.

Arista is proactively working with other leaders in the SDN space, co-authoring the NV-GRE standard with Microsoft, the VXLAN standard with VmWare and working with the OpenStack, OpenFlow and the Open vSwitch communities. The pillars of Arista's Software Defined Cloud Networking Vision are:

Pillar 1: Advanced Multi-path Cloud Topology: Scaling Cloud networking across multiple chassis via MLAG (Multi-chassis Link Aggregation Groups) at L2 or ECMP (Equal Cost Multi-Pathing) at L3 is a standards-based and scalable approach for uncompromised cloud networking. This insures effective use of all available bandwidth in a non-blocking mode, while providing failover and resiliency when any individual platform or link has an outage condition. MLAG and ECMP comprehensively cover all of the important multi-path deployment scenarios in a practical manner while not introducing any proprietary inventions. These technologies currently scale to 50,000+ compute and storage nodes, both physical and virtual, and heading towards 100,000+ nodes in the future. This scalability meets or exceeds 99% of cloud driven data centers. With the advent of next generation multi-core server CPUs, dense virtual machines and storage these types of uncompromised leaf-spine topologies with non-subscribed capacity, uplink, downlink and peer ports becomes paramount. The introduction of tunneling technologies as an overlay, leverages MLAG and ECMP without having to redesign these sophisticated and well proven traffic balancing, failover, and recovery protocols.

Arista's open platform approach and software modularity for allowing external systems to control the flows at edges makes hybrid L2/L3 topologies possible. These specific functions are best optimized in hardware for insuring wire rate performance, and augmented with Arista's Extensible Operating System (EOS™), a modern open software system designed for advanced network operations.

Pillar 2: Cloud Control: To effectively integrate with third party, open standards controller technologies, and purpose built controllers, the networking control plane (i.e. network operating system) must be standards based, open and modular. Many L2/L3 switching and routing vendors are trying to recreate decades of networking control plane architecture work with a non-modular, non-database centric, proprietary based (highly customized version of Linux) network operating system in response to Software Defined Networking initiatives. These operating systems are multi-year (or in several cases multi-decade) expensive undertakings with a history of vendor lock-in. SDN protocols (OpenFlow, OpenStack/Quantum, Open vSwitch), in combination with the L2/L3 IETF/IEEE traffic engineering specifications are the open approach for providing software programmable control planes in the future. Arista's Zero Touch Provisioning (ZTP) for automating deployments of compute racks as well as Latency Analyzer (LANZ) for detecting application-induced congestion and Advanced Event Management (AEM) are examples of innovative operational control.

Pillar 3: Network-wide Virtualization: By decoupling "the physical infrastructure" from applications, network-wide virtualization expands the ability to fully optimize and amortize compute and storage resources with bigger mobility domains and resource pools. It therefore makes sense to provision the entire network with carefully defined segmentation and security to seamlessly handle any application anywhere on the network (yet maintain tight security control and end user reachability). This drives economy of scale for cloud operators. This network-wide virtualization is an ideal use case in which an external controller abstracts the virtual machine from the network and defines the mobility and optimization policies with a greater degree of network flexibility than is currently available today. This requires both a tunneling/encapsulation approach and external APIs in which external controllers can define the forwarding path. The net benefit is much larger mobility domains across the network which is a key requirement for better optimizing resources (better return on assets). There are several tunneling technologies in development within the market today including VXLAN from VMware, NV-GRE with Microsoft, and others to be determined as part of the Open vSwitch developments. Arista has extensible architecture in which all of these can be implemented.

Pillar 4: Single Point of Management: Single point of management, which is often referred to as single interface or IP address, in which a network management system, including a controller, can program multiple switches concurrently is available from several vendors. A majority of these implementations are proprietary and often require “stacking” technologies. These stacking technologies consume ports and are closed; as result these approaches are costly and difficult to integrate with. The more modern day approach is to leverage a message bus technology such as XMPP, and RESTful based API’s with external management and controller systems. This approach is far more open, standards based and scalable. Specifically, a single interface for communicating to N number of switches can be layered on top of the traditional control plane and data path of a cloud network, without requiring a proprietary stacking technology. Arista’s CloudVision™ is a standards-based example of this, using XMPP messaging methods. This interface can be easily modified for future APIs based on OpenStack and/or OpenFlow specification updates.

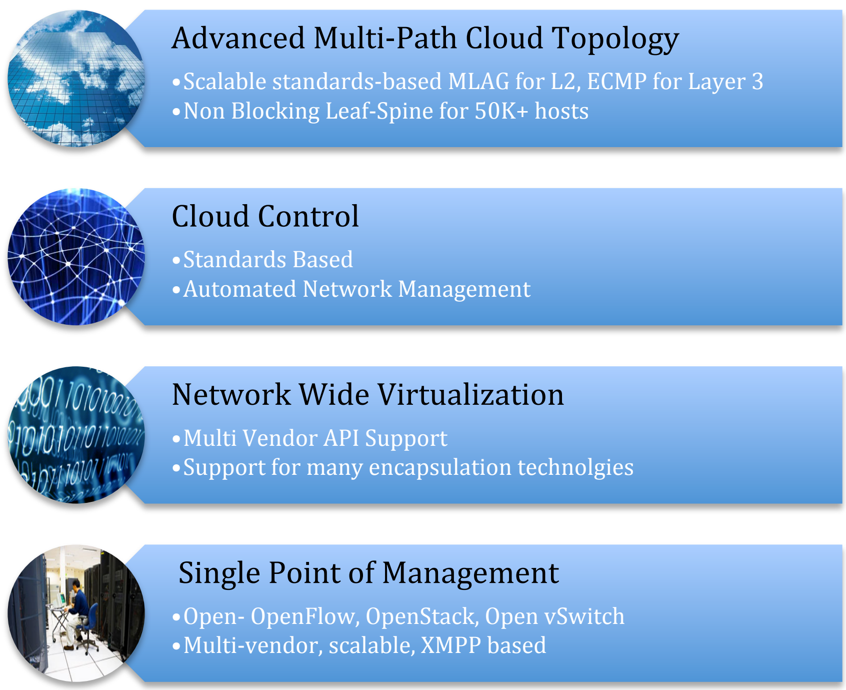

Figure 1: Four Pillars of SDCN

SDCN Applicability within the Data Center

The world in which an application binds to a physical address within the network is rapidly evolving as customers adopt cloud-hosting models, both internally, and with service providers. The ability to easily move applications as enabled with server virtualization technologies is a pervasive cloud requirement as this offers better capacity, failover, proximity placement, and security benefits. Cloud providers depend on these benefits in offering a profitable, secure, reliable cloud to their subscribers. Without these capabilities their implementations will not measure up to the service level expectations of their customers.

The complexities of moving applications are significant, as it requires real time changes within the server, network, storage, and security configurations. Prior to these new virtualized cloud inspired data centers all of these configurations were "hard wired" when data centers were designed around physical device constructs. While there are individual vendor efforts, semi custom script, workflow and run-book attempts to overcome the limitations of physical devices, these fall short for many customers as these are proprietary, non standards based, are difficult to maintain, and do not function at scale when there are thousands of subscribers. The ability to provision secure and segmented logical networks via software is critical to offering a virtual overlay.

It is relatively easy to drive scripts centrally when there are few subscribers, limited services (VLAN partitions), and infrequent changes. However, one of the underlying principals of clouds is to consolidate to a smaller number of very large data centers, with unprecedented asset utilization via virtualization. Additionally, to allow the subscribers to choose from a menu of options on-line including network security, addressing, quality of service and mobility services and to auto configure and modify real time. The largest cloud providers are looking for cloud solutions that scale to thousands of customers and hundreds of thousands of virtual machines, with change requests by the second. A centralized scripting and run-book approach for digesting provisioning requests from cloud portals, and issuing commands to the underlying infrastructure will not scale. The industry needs a more repeatable, scalable approach. Arista delivers the only network software for physical switches that is purpose built and capable of delivering software defined networking since introduction of its products in 2008.

Arista EOS: SDN Extensibility coupled with a Modern Operating System Matters

Traditional network operating systems are generally monolithic architectures. This means that all features reside in the same code tree, with little hierarchy or modularity. While often quicker initially to get to market, if a feature has a problem or if a bug fix is made to one of these features, in can affect the stability of the entire image. In contrast Arista EOS is a modular SW architecture running on a protected unmodified Linux kernel. Each EOS software feature, (or process) sits in its own protected address space. EOS modules can be easily modified, patched, and customized by customers without impacting other EOS modules. This offer unprecedented flexibility and integration opportunities with SDN based controllers.

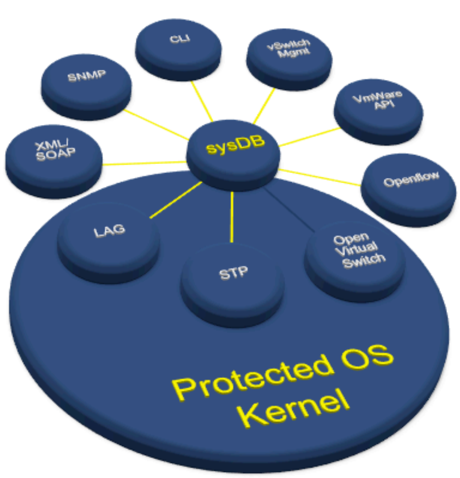

At the heart of Arista EOS is a state-sharing SysDB; a database that holds all session state and inter-process communications. For the application of SDN protocols or clients this separation of state from process allows the introduction of the SDN client without affecting the stability of the operating system. In addition the SDN client has access to the entire suite of services available on the switch (e.g. port mirroring, monitoring and health statistics, etc.). As a result, unique among network operating systems, Arista's EOS architecture delivers SDCN capability without affecting the stability of the network operating system while still retaining all of the EOS value added features and functions.

Figure 2: Arista EOS enables Software Cloud Defined Networking (SDCN)

Arista EOS enables SDCN Applications

The following block diagram offers a layered approach for understanding many of the technologies addressed in this paper.

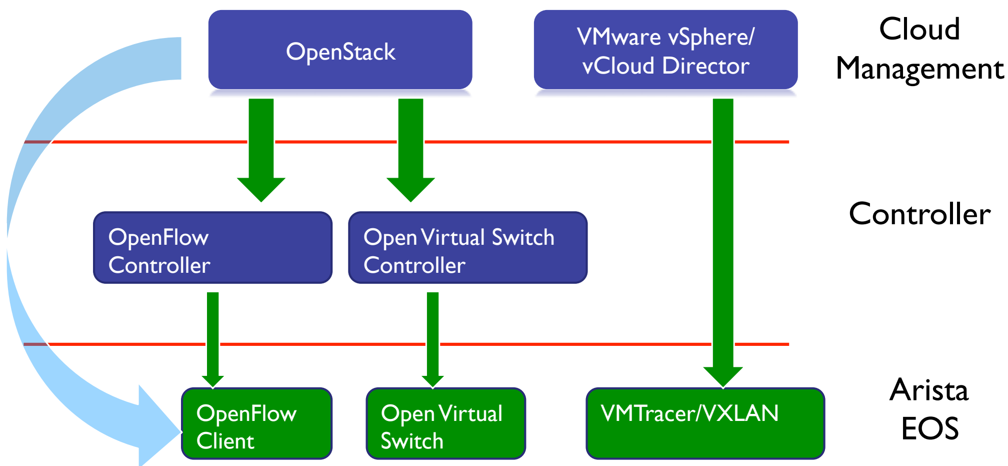

Figure 3: SDCN Layered Diagram

The open architecture of Arista EOS offers integration directly to an SDN controller (OpenFlow), or to a cloud management system (OpenStack). Both of these open source technologies maintain within their meta-models a management centric topology understanding. This management centric topology view defines the subscriber services available, that when requested through a cloud portal, can-be configured automatically. Depending upon the business requirements and data center management designs either management topology approach can be represented on top of an Arista network, with cloud compliant, industry standard API's.

The following paragraphs will identify how many of the SDN technologies work with Arista EOS.

Arista EOS & OpenFlow

The following provides a high level definition of OpenFlow

- OpenFlow is an industry specification for an open protocol to program flow tables in networking devices
- An OpenFlow controller is a device that uses the OpenFlow protocol to program networking devices.
- An OpenFlow controller has a two-way communication interface with the networking devices that make up the physical topology of the network.
- An OpenFlow controller can be based on multi-devices to insure scalability, failover and redundancy
- An OpenFlow controller can be X86 based, or virtual machine based, thus allowing customers to choose where to host based upon their performance, memory and design requirements.

OpenFlow controllers can eliminate the need of programming each networking device independently. These controllers have intelligence regarding where the services need to be placed or removed based upon the communication flows between the applications and subscribers (not based upon MAC or IP address tables). The OpenFlow specification (currently at v1.1) defines the interfaces between external controllers and the networking devices as well as many of the external controller functions in terms of setting up flows based upon service profile definitions. Because of its modularity, scalability and open interfaces, Arista's EOS architecture is ideally suited for the insertion of the OpenFlow protocol. Like other Arista protocols, the Arista OpenFlow client will be standards based providing support for a variety of OpenFlow controllers, both open sourced and with vendor specific value added controllers.

Arista EOS and Open vSwitches (OVS)

Open vSwitch (OVS) is currently a virtual switch that runs in a hypervisor, and is the default switch in XenServer 6.0. XenServer is open source code and is used by Citrix, and Red Hat for Citrix XenServer and Red Hat KVM respectively. While XenServer running on an X86 based COTS server is the primary use case today for OVS, future OVS applications will make use of hardware acceleration and have OVS running as a control plane protocol on network switches. As with with OpenFlow, Arista EOS delivers a unique networking software platform for the introduction of an OVS onto a network operating system. Arista can currently demonstrate OVS integration, with EOS software integration. Future Arista switching chassis will support OVS in hardware, for offering best in class hardware acceleration.

Arista EOS and OpenStack / Quantum

Arising out of work done within NASA, OpenStack is an open-source project that aims to deliver a scalable cloud operating system, providing Open Source software that maps data center elements such as servers, networks, and data storage to specific applications and services. OpenStack is comprised of five elements:

- Nova: provisions compute services, akin to Amazon's EC2 service
- Swift: object oriented storage, similar to Amazon's S3 service
- Glance: a virtual machine image management service

- Horizon*:* provides cloud-based resources through a self-service portal
- Keystone**:** A common authentication system for multiple log-in credentials, including username/password, token-based and AWS-style logins.

Initial implementations of OpenStack network provisioning were done via the Nova working group, primarily focusing on configuring VLANs. Recently the OpenStack community has started an incubation project called “Quantum”. Quantum is designed to provide network connectivity as a service to other devices managed by other OpenStack entities (e.g. Nova). Quantum itself is made up of two elements: the service itself, and a plug-in (typically vendor or technology specific). The Quantum service handles managing network definitions, and authentication.

Today, Quantum provides an API for the management of network segments, and an API for specific vendor plug-ins. The vendor plug-in owns every action necessary to map the Quantum abstractions to the physical networking devices. Quantum is being developed with a layered approach. Network administrators can interface directly with Quantum and request changes that program the networking devices directly, or they can interface through Quantum, and make requests to an SDN controller which in turn communicate to both the network devices, and the virtual switches. Arista will support Quantum plug-ins for enabling both approaches.

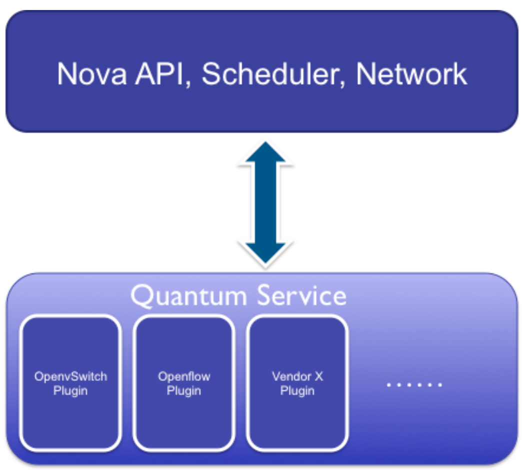

Figure 4: Quantum Architecture

Arista Implementation in OpenStack/Quantum

Customers interested in using Quantum as the networking provisioning element in their cloud networking infrastructures can utilize Arista switches by either using the OpenStack Nova project today to provide simple VLAN provisioning and in the future have the following capabilities:

- Utilize a Quantum OpenFlow Controller or OpenFlow Controller plug-in and an Arista OpenFlow Client
- Utilize the Quantum Open vSwitch or Open vSwitchV Controller plugin and an Arista Open vSwitch Client

Arista currently has support for OpenStack with partners that provide support for Arista with the OpenStack Nova project or in the future Quantum support.

Arista and VMware's vCloud Director

Arista currently provides support for VMware environments with VMTracer, a capability that links Arista switches to VMware's vCenter as shown in Figure 5. VMware's vCloud and vSphere's/vCenter are the best example of a widely deployed controller in data centers. Arista VMTracer coupled with vCloud Director delivers an adaptive infrastructure that automatically responds to vCloud Director virtual machine management changes and configures the network in response to the service requests at the portal layer. VMTracer utilizes the published vCloud Director APIs and works across all editions of vSphere. This offers unprecedented visibility into the virtualized environment, seamless integration into a familiar industry-standard CLI, and automatic configuration of service requests with native VMware vCenter integration.

VXLAN: Ideal Extension of VLANs across L2 and L3 boundaries

VMware and Arista co-authored an IETF draft specification called VXLAN that is designed to enable a virtual machine to be deployed on any server regardless of the IP subnet the physical ESX host is in. VXLAN is an encapsulation mechanism that runs between virtual switches and enables virtual machines to be deployed and moved on or between any servers attached to the network. This enables the IT department to deliver a scalable network architecture that supports capacity on demand and workload mobility regardless of geography and IP addressing.

Arista will introduce VXLAN Extensions to VMTracer inline with the commercial availability of VXLAN from VMware controllers. These extensions will provide for the automatic configuration of VXLAN Virtual Network Identifier addressing and IP multicast groups and provide support for a hardware VXLAN gateway for those devices that do not have native VXLAN support such as storage arrays, load balancers, or firewalls.

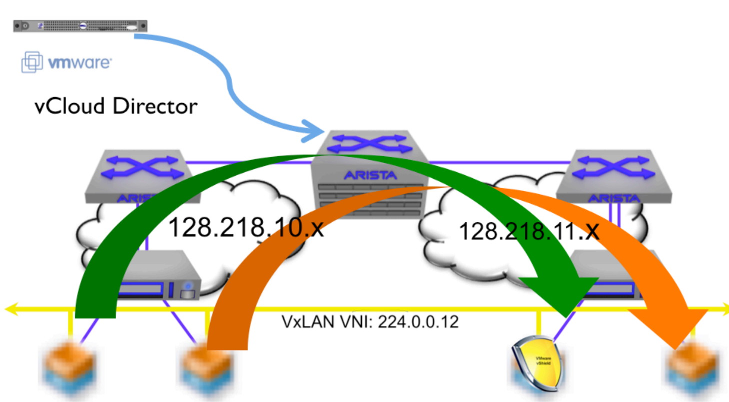

Figure 5: vCloud Director programming VXLAN tunnels to enable vMotion across L3 boundaries

Summary: Arista EOS makes SDN a Reality

With the multitude of SDN choices available to IT organizations, choosing the right SDN technology for their data center and business requirements can be a challenging exercise. In many cases one-architecture may not be enough as the requirements may dictate different architectural deployments for different business operations or business units. Laying down a foundational network architecture that provides SDN operational flexibility as well as high performance protects IT networking investments. IT organizations that standardize on Arista EOS are ensured that they can deploy and choose the correct SDN based architecture for their business needs.

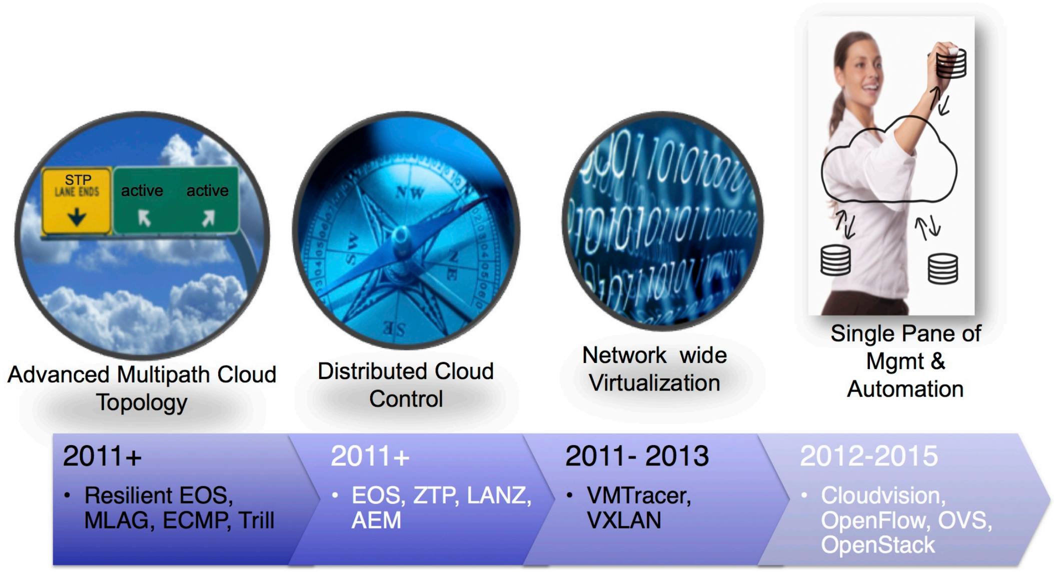

Figure 6: Evolution of Software Defined Cloud Networking

For more information, please go to www.aristanetworks.com

Arista Networks
5470 Great America Parkway,
Santa Clara, CA 95054
Phone: 408-547-5500
Fax: 408-538-8920