Cloud Grade Networking for Live Production

The growing demand and ease of consumption for content continues to skyrocket and the Media industry continues to reinvent itself in response. The ongoing rapid adoption of commodity infrastructure, cloud technologies, and remote production is driving transformation in the media production, contribution and distribution industry.

The adoption of IP based infrastructures is key to the success of the business transformations that are underway in the media industry and content providers are using the scale of IP based distribution to differentiate their offerings, become more agile, and reach wider audiences.

Producers are re-imagining their workflows to maximize resource-sharing and enable distributed production, delivering higher production values whilst reducing idle time on technology resources and the amount of time and money expended on travel.

Given the global audience and competitive landscape, content providers and producers are migrating their creation, production and distribution workloads to IP-based infrastructures to deliver a more productive, flexible, streamlined, and cost effective solution.

Introduction

Media production companies, be they pre- or post- production, or real-time broadcast, are dealing with an explosion of content as they race their productions to market. Digital capture and editing technologies, along with next-generation animation workflows, create multiple image streams that must be passed among realtime and non-linear broadcast workflows, and then to IP destinations in varying formats.

CGI and animation production studios were the first to leverage price/performance gains and investment protection of mainstream IP based infrastructures. Image, file and transport protocols streamlined file based workflows, allowing the use of COTS compute platforms running over standard IP networks. These open architectures simplify scaling and technology evolution, and have allowed much of this infrastructure to continue its transformation from on-prem to in-cloud.

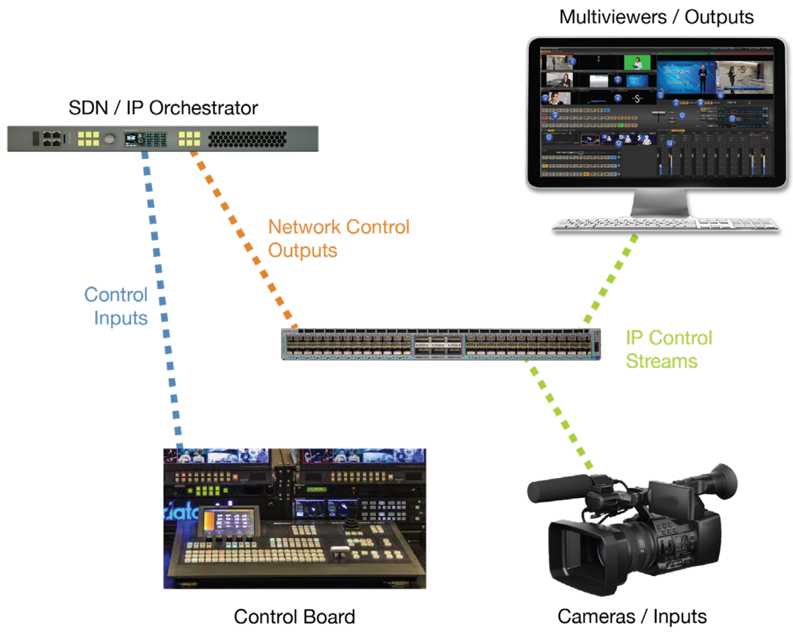

The SDI to IP challenge

Migrating real-time broadcast workflows from SDI to IP has been a challenge. The high bandwidth, low latency, real-time nature of Live Production brings significant technological hurdles. Additional timing and control standards are needed to facilitate interoperability and scaling whilst ensuring synchronization and quality. Equally important is the development of open, standard orchestration APIs that quickly and reliably translate the operators actions at the video mixer or camera shading position to the affected traffic flows in the IP switching infrastructure.

The industry's standards bodies, such as SMPTE, AMWA, JT-NM, AES and the VSF have created the standards necessary for the successful implementation of the move from SDI to IP based workflows.

Commercial solutions, based on these standards are now widely available and demonstrably interoperable, and when combined with suitable network infrastructure, are enabling new forms of content creation while supporting familiar workflows; all whilst disrupting the traditional cost model. The IP enabled Live Production workflow increases flexibility, scale, quality and reuse of content while reducing production and delivery costs.

Broadcast Requirements Overview

Critical to the reliable delivery of Live Production workflows across an IP infrastructure is network switching bandwidth, control and feedback, timing synchronization, and telemetry. The utility and flexibility of network switches is measured by the breadth of capabilities it can offer to applications used by content creators.

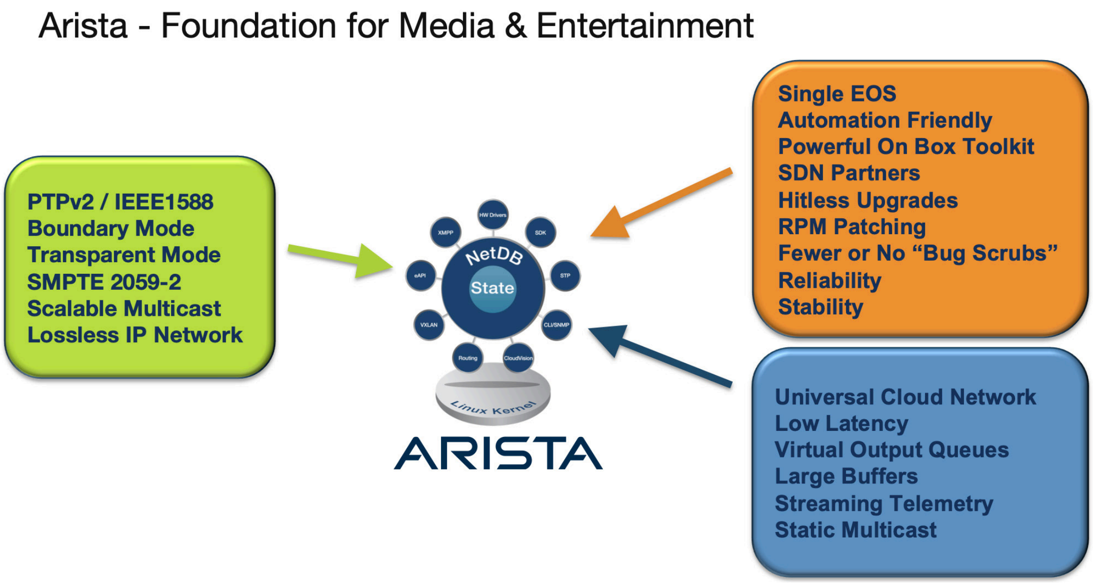

IP - Transforming content creation

- Enabling flexible and agile workflows, resulting in efficient resource sharing and distributed production models
- Increases capacity to support current and future formats, enabling more visually immersive theatrical workloads and broadcast workstreams
- Converges content creation, rights management, transcoding and distribution
- Delivers an open, standards-based transport, supporting economical, and commodity based authoring, rendering, transcoding and storage system
- Provides native connectivity between on-prem and in-cloud environments

Bandwidth:

Whether broadcast, animation of post-production, higher imaging formats evolving from SD to 2K, 4K, and now 8K resolutions, exponentially grow file sizes and broadcast stream bandwidths.

While consumer directed, low frame-rate, H264/5 compressed HD video streams use megabits of bandwidth; production quality, uncompressed HD and 4K streams demand multi-gigabit data rates.

Depending on format, IP based uncompressed video switching systems require 10Gbps or greater speed interfaces to support the demands of flow hungry devices like multiviewers and video mixers.

Fortunately, cloud grade data center networking switches support these data rates in compact, high density configurations. Wire speed ethernet switches can provision 128 X 25Gbps ethernet connections per rack unit (RU) in a power efficient enclosure. These platforms are wirespeed/non-blocking and also support 100 or 400Gbps ethernet interfaces providing uplink/tie-line functions for distributed broadcast workflows.

Many broadcast vendors are creating devices with one or more 100Gbps interfaces to allow for uncompressed UHD signals to be transmitted and received in bulk, thus establishing a need for even greater bandwidth between nodes. Large production facilities can use 400Gbps based infrastructure to achieve the necessary scale and density today, with the future bringing 800Gbps.

In a streaming multicast environment, packet drops are unacceptable, and with the unavailability of TCP/IP retransmission, once a packet is gone, it's gone. SMPTE ST2022-7 hitless merge techniques, and application layer FEC (Forward Error Correction) can be used to mitigate the risk.

Deep buffer switches are better at avoiding packet drops and can deliver hundreds of times more packet buffering capability compared to shallow buffer systems. These extra buffers are used to smooth out the micro-bursts that occur as a result of the instantaneous high bandwidths experienced while processing the interface speed changes, or "zipping" together multiple flows at an "in-cast" point like a vision mixer or multiviewer.

Live Production workflows require low latency, and so it is often assumed that buffering in the network infrastructure is detrimental. However, we are dealing with an environment where the majority of traffic is multicast; one or many to any number from 0 to all! This multicast model works well for these types of workloads. For example, a camera feed is often needed at the vision mixer, shader position, ISO recorder and multiple multiviewer locations. These atypical flow patterns often result in 'in-cast' situations, where many flows converge to egress on the same interface such as flows to a video mixer, or multiviewer. At this in-cast point in the network, many asynchronous, large and potentially bursty flows need to be queued temporarily in order that they can all be 'zipped' back together into a single flow at the egress port. For a multiviewer or vision mixer, this could be 60+ high bandwidth flows. The queueing requires buffers and sophisticated scheduling algorithms that can minimise buffering needs, but ultimately the need for sufficient buffering cannot be avoided in these high bandwidth, multicast environments.

Control and Feedback:

While both file and streaming workloads benefit from increased bandwidth, broadcast workflows also require network control and feedback mechanisms to ensure reliable operation and broadcast quality continuity.

Amongst the key requirements in a high performance IP based Live Production network are network based real time replication services, like multicast, plus traffic steering APIs, so that video and audio feeds can be controlled by software based video orchestration systems. IP multicast services simplify duplication of real time data streams to multiple consumers. Features like IGMP snooping and Protocol Independent Multicast (PIM) routing allow the network to control broadcast streams, maximizing network resource utilization while minimizing the burden and cost of network connected cameras, playback devices, multiviewers, etc. Fast and reliable IP multicast joins and leaves replicate the behaviors of SDI routers. Inconsistent or slow multicast behavior compromises application performance and risks making the Live Production environment unresponsive.

Another key to optimizing Live Production applications are the programming interfaces (APIs) that help connect user actions on vision mixer or router control panels to network traffic actions on the underlay IP network. SDN (Software Defined Networking) orchestrators are the middleware that relay actions executed on a video producers' control panel to the network operations that fulfill those actions. Video source joins, merges, patches, and other actions are orchestrated at the network layer, synchronizing the networking action of specific devices and the networking hardware itself. Broadcast controllers also provide necessary APIs to receive and send control messages to automation systems, control panels, cameras, playback devices, video mixers and other devices in the broadcast production network. The reliability and confidence of a broadcast workstream depends on the quick and faithful execution of the underlying networking infrastructure, their APIs and their interaction with the Broadcast Controller. The producer expects a real time "feel" from their control button push. Anything less is unacceptable.

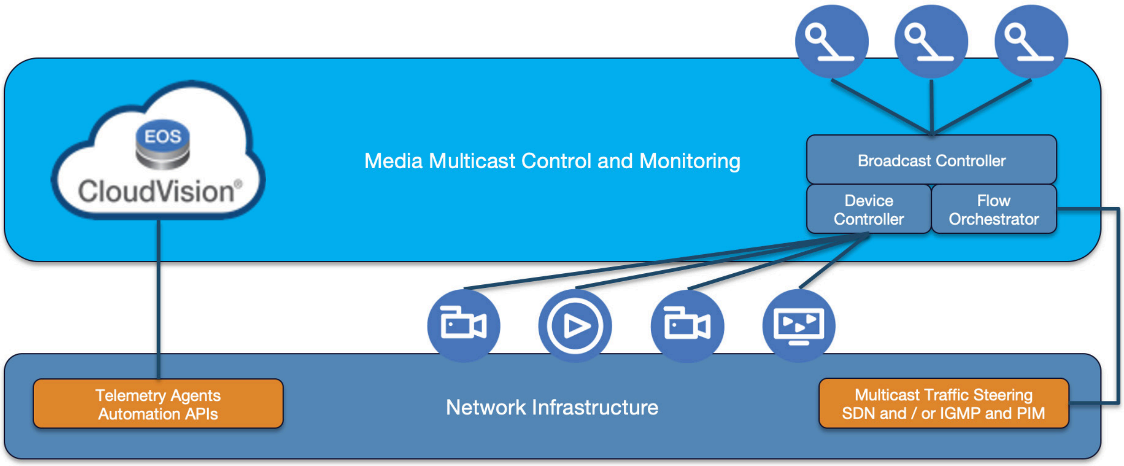

SDN (Software Defined Networking), or Flow Orchestration, is a broad term that describes the process of taking active control over exactly where the multicast flows travel in a network. To be most effective, it needs to be both bandwidth and path aware.

Bandwidth management is a key component of any multicast network carrying uncompressed SMPTE ST2110 flows over a complex multi-switch topology. Flows need to be actively load balanced over the infrastructure to maximize throughput. The paths that these flows take need to be carefully managed to ensure diversity for SMPTE ST2022-7 flow pairs.

Adding SDN to a Live Production infrastructure also provides the benefits of determinism in the case of component failure, or component maintenance. The SDN system can be used to provide re-route capabilities that are bandwidth and path diversity aware.

Utilizing the open standard, Openconfig protocol, broadcast vendors can integrate their control systems directly with the Arista infrastructure providing fast multicast switching speeds. When implemented correctly, the data network will not have to rely on PIM and other multicast routing protocols.

Timing and Synchronization:

Platforms used in Live Production workflows must be tightly synchronized to ensure faultless continuity of real time video/audio programming. Video broadcasts of sporting events, ranging from photo finishes imagery to multi angle images used by replay officials, help underscore the importance of timing synchronization. Synchronization leads to lower latencies and end-point costs, but perhaps most importantly, ensures lipsync and multi-channel audio coherence.

SMPTE ST2059-2 has adopted the IEEE-1588 precision time protocol as the service for time synchronization in IP based broadcast networks. PTP is IP based and provides services for ensuring accurate time, frequency and phase distribution between grandmaster, and device end-points. PTP sources, typically GPS locked grandmaster clocks, provide timing information to clients. The PTP protocol provides mechanisms by which both the time offset from master, and the network delay between master and end-point can be determined. This leads to the ability to provide very accurate frequency and phase synchronization across the network. IP based audio installations have long used PTP to synchronise end-points, and in these environments, PTP aware infrastructure was not always needed. However in larger studio networks, with the higher accuracies that SMPTE ST2059-2 requires, PTP distribution does not scale well in non-PTP aware infrastructure, so it becomes necessary to introduce Boundary or Transparent Clock capable network switching, to ensure reliable, accurate PTP delivery.

As packets travel through a network, they experience latency in each switch hop. This latency is caused by the packet being stored in buffers, before being scheduled to be sent to the next hop. The latency will depend on how much traffic is present in the network, and whether the packet has been prioritised. Each packet experiences a different latency - this is called jitter. As you can imagine, it becomes more difficult for an end-point to accurately determine the time from PTP messages if these messages are subject to unknown latency and jitter. PTP aware switches use mechanisms called Boundary or Transparent clocks, to eliminate the packet delivery jitter of PTP packets - thereby making it possible to achieve the PTP accuracies that are needed for SMPTE ST2110 based Live Production facilities.

Network switches that support PTP Boundary mode help expand the capacity of the broadcast network by distributing the time synchronization workload. Switches supporting PTP Boundary mode, synchronize with an upstream master and provide accurate PTP information to downstream devices.

In contrast, transparent mode PTP time keepers act as a helper, adding individual switch occupancy info to each PTP event packet, allowing the end-point to dial out the jitter. Transparent clocks don't provide the offload service that Boundary clocks do, so the grandmaster still communicates directly with each end-point.

Distributed time synchronization using Boundary mode offers resiliency and security advantages. Boundary switches continue providing timing data even if a master clock becomes incapacitated. Once the master clock has recovered or a backup has taken over, the network recalibrates without losing a step, thus ensuring consistent, reliable broadcast operations. Boundary clocks can also be configured to prevent rogue or badly configured end-points from assuming the grandmaster role. Boundary clocks can provide very useful telemetry to help verify PTP performance and troubleshoot in the case of PTP problems.

Telemetry

Monitoring network health and performance is important for any workload, but especially so for Live Production. For decades, SNMP has been the mainstay protocol for providing networking telemetry to operators. Unfortunately, while SNMP provides a solid framework for information retrieval, it lacks the ability to provide the rich and timely information that is needed to monitor a high performance media network.

Arista EOS is built with a state based, "publish and subscribe" database at its heart. Arista leverages this modern architecture to stream every state change from every device into a central, cloud-based view, or to 3rd party monitoring solutions.

The benefits in a Live Production media network is that monitoring tools and SDN orchestrators can utilize this near real-time streamed information to provide enhanced monitoring, alarming and control feedback to maximise performance and visibility, allowing Broadcast Engineers to focus on their broadcast workflows instead of their network.

Arista Networks: Preferred Choice of Media and Entertainment Studios

Arista Networks is the leader in the cloud data center networking market, with thousands of customers who have deployed Arista data center-class switches for their mission critical workloads. Many of these applications require hundreds or thousands of server and compute nodes to: process millions of messaging streams within microseconds; stream hundreds of real-time broadcast, video, and movie streams across the internet; or reliably move large data files between data centers at wire speed.

Arista has been successfully deployed and is trusted by many broadcast and media networks around the world to provide reliable and transparent infrastructure for some of the largest sporting and cultural live events globally.

Arista offers a complete product portfolio supporting both private and public cloud. Arista's portfolio of wire-speed, low latency, 1/10/25/40/50/100Gbps and 400Gbps Ethernet switches use the Extensible Operating System (EOS®) that is common to the product family. This operating system is modular, extensible, open and standards-based, supporting fast turn-up and extensibility of sophisticated and customizable switching features for studio production data centers. EOS is uniquely customer centric: administrators can create their own scripts and utilities for specialized network configurations and operations.

The power and modularity of EOS also lends itself to serving public clouds in the role of an aggregation virtual router. As cost and time to market needs dictate, content owners will deploy virtualized rendering or post production workloads across a variety of public clouds. The EOS vRouter provides network administrators a standardized virtual routing platform that's identical to, and can be managed like all other Arista platforms running EOS. EOS vRouter not only supports the same configuration management and telemetry available in Arista switches, but also provides CloudTracer telemetry that delivers near real-time reporting of cloud network and application response time delay, throughput and reliability.

The vRouter is available on all popular cloud platform services and delivers competitive price/performance. vRouter and its siblings, vEOS and container cEOS, can also be used for modeling network designs in a fully virtualized VMware, Docker containers, or KVM environment. Customers can leverage the same network OS to streamline administration and networking of private and cloud based data centers. This is just one example of how Arista redefines the process and tools used to design, deploy and maintain modern media production data centers.

EOS: Extensibility, Automation and Reliability

Arista's EOS has a fundamentally unique architecture. Its foundation is a multi-process state sharing architecture that separates state information, protocol processing, application logic and packet forwarding. EOS system state and data is stored and maintained in a highly efficient, centralized System Database (SysDB). SysDB uses state of the art publish/subscribe/notify models and APIs. This architecturally distinct design greatly improves reliability, maintainability, and extensibility. All this ensures product quality with faster time-to-market for new features that customers require.

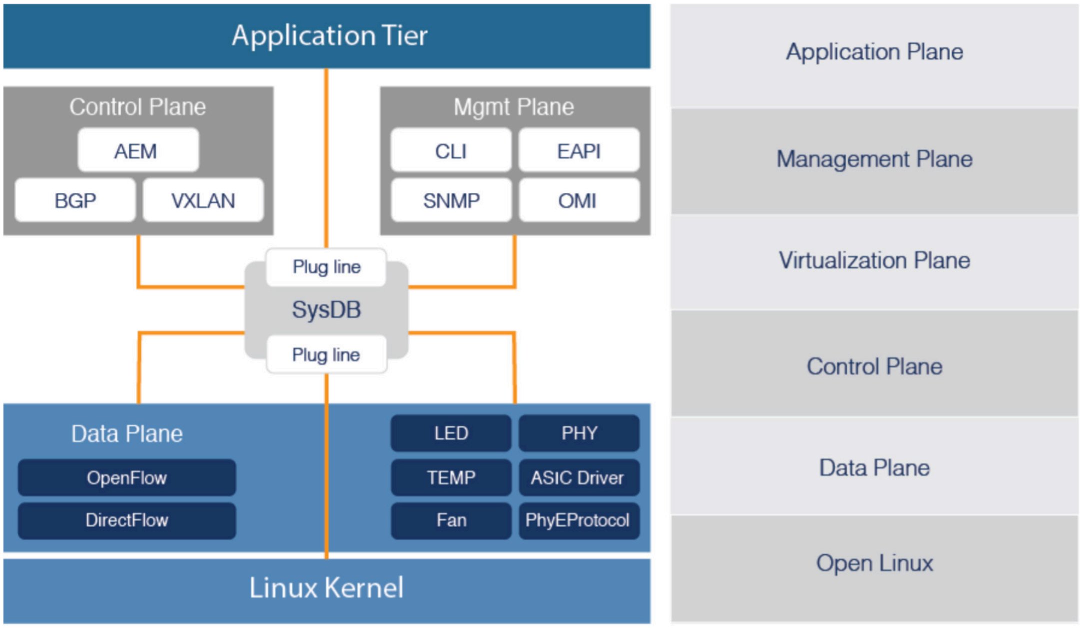

With leaf and spine architectures becoming the norm for large scale Live Production networks, IGMP and PIM are not the most effective method of distributing multicast groups across a network, as they do not provide bandwidth aware flow provisioning.

In a large Live Production environment, there may be tens of thousands of multicast flows, thousands of these could be 10Gbps each in a UHD workflow - and without bandwidth aware orchestration, traditional hash based load balancing can result in disastrous hotspots.

Arista's Media Control Service or MCS, is a Network Control layer that can provide the topology and bandwidth aware orchestration services that overcome the limitation of IGMP and PIM as outlined above. A simple northbound API allows for easy integration with IP aware Broadcast Controllers, and tight integration with the underlying Arista infrastructure providing maximum visibility of the network state back to the Broadcast Control layer.

The Big Picture: Media Architecture

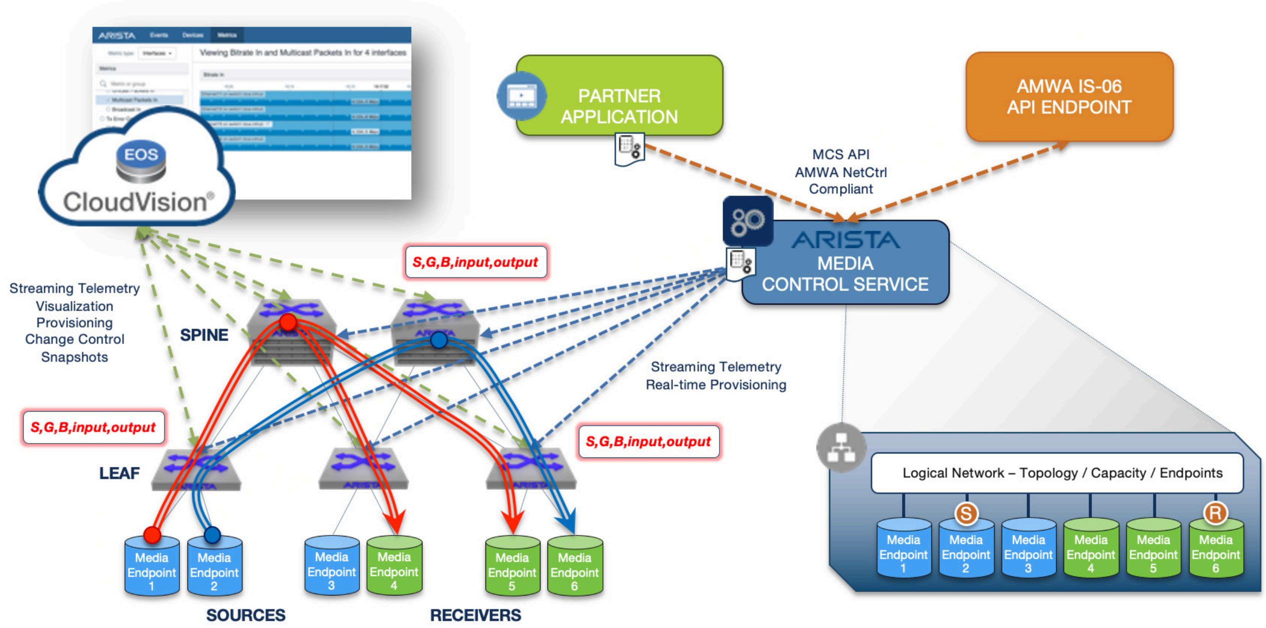

EOS has been built from day one for automated change management: to allow for robust integration with technology ecosystem partners and network programmability at cloud-scale. The Arista EOS architecture allows full programmatic access to all aspects of the software, including the event-driven, state-based infrastructure known as SysDB. Already proven at cloud scale, the EOS approach provides automation and integration flexibility to help gain the efficiencies and cost-savings leveraged in DevOps environments.

With Open APIs and developer contributions available in public repositories like GitHub, administrators can create or adapt their own DevOps tools, addressing a wide range of operational needs such as provisioning, change management and monitoring. Administrators can implement solutions based on tools such as Puppet®, Ansible® or Chef®, and can integrate Arista provisioning tools such as ZTP Server.

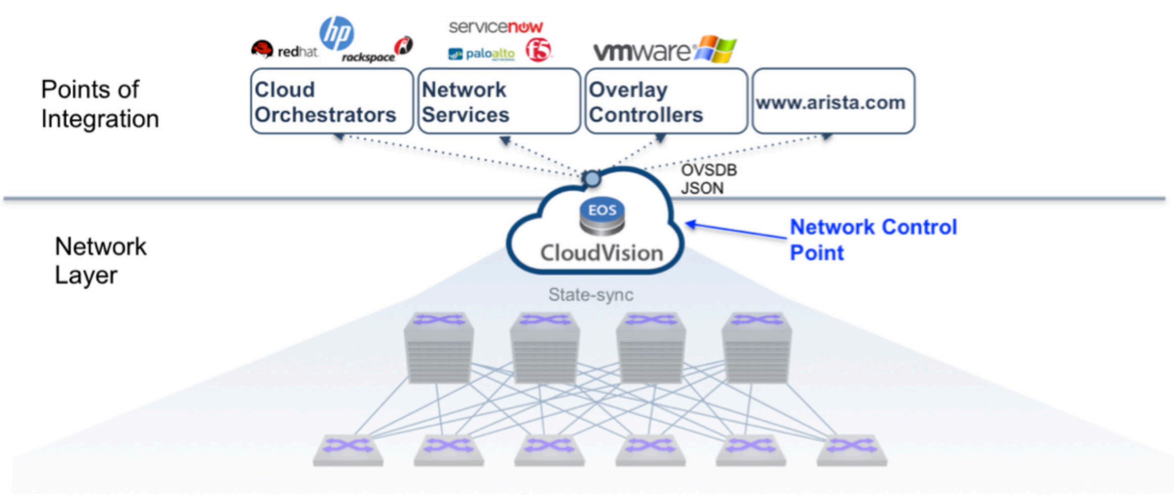

For administrators who have neither the time nor personnel to develop administration systems, Arista has developed the CloudVision® suite of tools for network system administration. Available in an easy to deploy virtualized platform, CloudVision can address audit and compliance requirements, including device inventory, configuration, and image version management. CloudVision provides the means to automate software life cycle management without service impact by using Smart System Upgrade (SSU) and ECMP maintenance mode mechanisms for faultless, and impact free management. Cloudvision also provides near real time telemetry monitoring, including multicast monitoring for broadcast workstreams.

Consistent with Arista's open and programmable model of coexisting in a robust management ecosystem, CloudVision is equipped with open and complete APIs for integration with an organization's existing management and monitoring framework.

Arista 7000 Switching Family: Platform Architecture Built to Serve Any Size Data Center

Arista's portfolio offers many features that improve the performance and reliability of broadcast and media post-production pipeline, simplifying migration of broadcast, CGI and computer animation applications to the data center. Regardless of workload or scale, there's an arista platform running EOS that fits the workload. Administrators can deploy low latency 1GBASE-T switches for broadcast work streams or deploy deep buffer switches for lossless large-scale workflow transfers. Administrators can utilize installed RJ-45 connected copper cabling for 1Gbps workstation and legacy server connectivity in the data center. In most cases, the same UTP infrastructure can also be used to deploy 10GBASE-T networking on workstations or servers sporting new 10GbE UTP adapters. This flexibility saves time and money when retrofitting workstation seats or server racks for next generation production workloads.

Similar price performance efficiencies exist for data center architectures that need to support increased bandwidth 4K workloads. Industry standard 25GbE uses the same grade twinax cable and fiber optics used for current 10Gb Ethernet.

25Gb server adapters and switch connections cost the same as 10Gb, so managers can now deliver more than double the network bandwidth for virtually the same budget. These switching platforms also provide cost effective 40, 50, 100 and 400Gbps for high bandwidth applications such as storage and inter switch connectivity.

Media and Entertainment Architecture

As demand grows, administrators can efficiently scale their high-performance network by leveraging data center leaf-spine architectures: start with a single RU, high density, fixed top-of-rack switch then interconnect racks with an aggregation spine layer of switches. Scale and bandwidth are determined by the speed of network cross connects, the number of distribution layer switches and the data paths between leaf and spine layers. Arista offers a full line of switches that support this architecture.

For larger scale applications, Arista offers the modular 7500R2 series switching chassis with 4, 8, 12 and 16 slots. Each slot supports up to 9.6Tbps of bandwidth to support incredibly high density 10, 25, 40, 50, 100, and 400G Ethernet connectivity. The 7500 implements a resilient Virtual Output Queuing (VOQ) architecture that ensures near 100% fabric efficiency without traffic loss, fabric hot spots, or head of line blocking (HOLB). This critical feature ensures that workflows between hundreds of servers or real time broadcasters cannot be impacted by isolated network events or congestion to a particular device.

If the roadmap to 800Gbps is important, then the 7800R3 series switching chassis can be considered the next generation of the 7500R series, sharing the same benefits, but with a max throughput of 460Tbps.

Conclusion

The networking revolution in Live Production is opening new possibilities for scale, flexibility, distributed production and resource sharing driving up production values, and driving down costs. Standardization in platforms and tools improves production economics, opening these capabilities to a wider audience. Production pipelines are greatly streamlined by automation and enhancements that are artist and administrator friendly. This streamlining and automation must also apply to the studio production network infrastructure.

Arista Networks leads the industry with a software-defined, cloud-grade networking architecture, delivering leading price/performance, scaling, automation and administration. Driving Arista's best in class switching platforms is its innovative Extensible Operating System (EOS): a Linux based platform that supports automated healing, reconfiguration and extensibility. EOS is consistent across the entire product portfolio ensuring consistent reliability and scalability across all production use cases. To expedite provisioning, changes and upkeep in the growing data center while containing administrative costs users can leverage CloudVison along with Arista's MCS, or develop a bespoke administration system using industry standard APIs, and COTS Dev-Ops tools. These automation tools also avoid costly errors that may impact productivity.

With Arista, media production companies can take advantage of the same management; performance and scaling efficiencies realized by high-performance computing (HPC) and cloud service providers. Studio production resources will perform better, will be more reliable and cost less, so there's more time and resources for the studios' artists, designers, and producers.

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