

Cloud Grade Networking for Broadcast, Digital Media Creation, Post Production, Content Distribution and Storage

Developments in digital broadcast, VFX, and animation technologies are expanding the creative horizons of producers and talent. Concurrently, the growing demand and means of consumption for such content is astounding! The scale and complexity is measured by the number and types of devices connected to the Internet. Content providers are innovating and expanding IP based distribution solutions to differentiate their offerings and adapt to growing audience demands. Likewise, producers are identifying and creating new content using real time, multi-stream workflows made possible with multi-vendor production systems running on Ethernet-IP switching. Given the global scale of audience and competition, content providers and producers are migrating their creative, production and distribution workloads to IP based Ethernet infrastructures to deliver a more productive, streamlined, and cost effective solution.

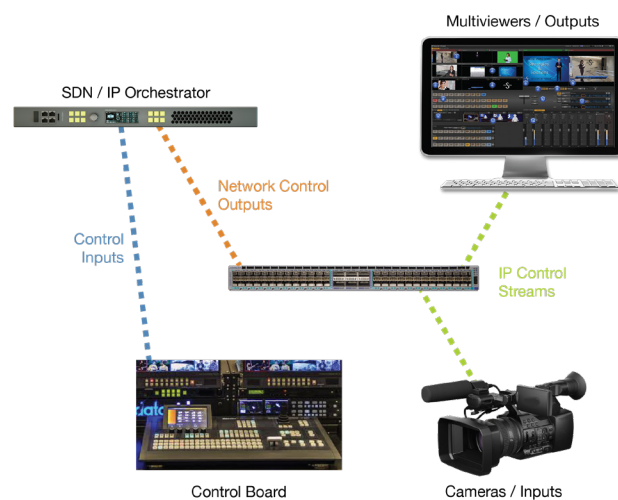
Accelerating Content Growth

- Increases capacity to handle 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

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 parsed among realtime and non-linear broadcast workflows, and then to IP destinations of varying formats.

CGI and animation production studios were the first to leverage price/performance gains and investment protection of mainstream IP based Ethernet infrastructures. Image, file and transport protocols streamlined file based workflows, allowing the use of COTS compute platforms running over standard IP Ethernet networks. These open architectures simplify scaling and technology evolution. For example, higher fidelity display technologies are economically fulfilled using more processing and network capacity for image rendering.



Migrating real-time broadcast workflows from SDI crosspoint switches to Ethernet based IP is a greater challenge. Additional timing and control standards are needed to facilitate interoperability and scaling while ensuring quality and synchronization. Equally important is the development of open, standard orchestration APIs that quickly and reliably translate the actions on a video control board to the affected traffic flows in the IP-Ethernet switching infrastructure. The standards being developed at SMPTE with the support of the Video Services Forum and other industry standards groups are now being applied to production ready IP Ethernet platforms. These solutions, available from industry stalwarts as well as upcoming innovators, are enabling new forms of content while supporting current and future generations of capture, editing and display formats; all while greatly minimizing future migration costs. The foundation of this transformation are high-speed, scalable, switched IP networking infrastructures that reliably and easily interconnect a universe of end point systems.

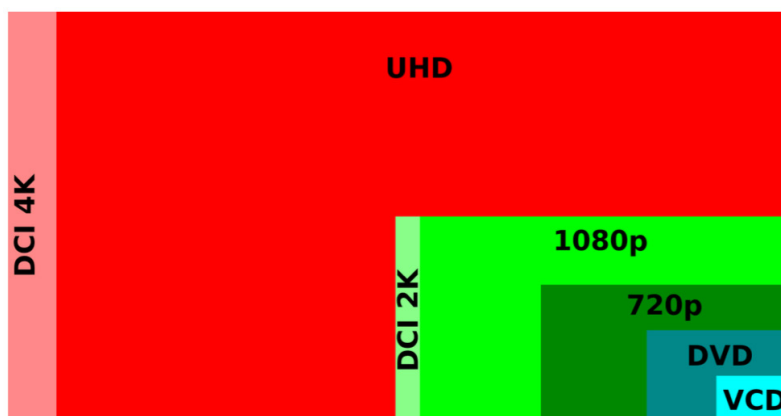


Requirements Overview

Whether content is stream or file based, network switching bandwidth, control and feedback, timing synchronization, and telemetry are the foundational requirements for reliable delivery of media across an IP network infrastructure. The utility and flexibility of network switches is measured by the breadth of capabilities it can offer to applications used by content creators.

Bandwidth

Whether broadcast, animation or post-production, higher imaging formats evolving from HD to 2K, 4K, and now 8K resolutions exponentially grow file sizes and broadcast streams.



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.

Table 1: Data rates of SMPTE 2110-20 compressed 2K and HD video applied to Ethernet

Scan format/fps	Gn/second	Ethernet Support	Simultaneous streams (full duplex) supported
2160p @ 59.94%	10,600	25G	2 x 2
1080p @ 59.94%	2,653	10G	3 x 3
1080p @ 29.94%	1,325	10G	6 x 6
720p @ 59.94%	1.178	10G	6 x 6

Depending on format, IP based uncompressed video switching systems require 10Gbps or greater speed interfaces to support the demands of multiviewers and compositors. Fortunately, cloud grade data center switches support these data rates and higher, in compact, high density configurations. Wire speed ethernet switches can provision 120 X 25G ethernet connections per rack unit (RU) in a power efficient enclosure! These platforms are wirespeed/non-blocking and also sport 100 Gbps Ethernet interfaces providing interconnect/tie-line functions for distributed broadcast workflows.

Among the new switching technologies are deep buffer switches which can better avoid data loss. In a streaming multicast environment, packet drops can prove costly since a lost Ethernet packet can invalidate an entire video frame. Deep buffer switches are better at avoiding packet drops than their shallow buffer counterparts. Deep buffer platforms deliver hundreds of times more packet buffering capability compared to white box switch on chip (SOC) systems. Induced congestion latency is minimal: 1.2 microseconds for 1500 byte frame over 10gb Ethernet. Less than six times that for a 9KB jumbo frame. Receivers have time to still use the frame and avoid looping stale video or worse!

When Buffering is Better than Dropping

An uncompressed 24 bit HD frame has a lot of data: 6.22MB! Translated into 4100 ethernet packets, they traverse a 10gb link in 50.3 micro-seconds. However, if you add from one to five microsecond switching processing latency per packet, the network transmit latency can grow to 200ms. IP broadcast platforms are built to buffer video frames to accommodate expected latency but if a network packet is lost due to a shallow buffer switch, the entire HD frame is lost and It's pointless to retransmit a stale frame. Buffering in the network helps avoid retransmissions due to congestion. It's a worthwhile investment for broadcast critical workloads.

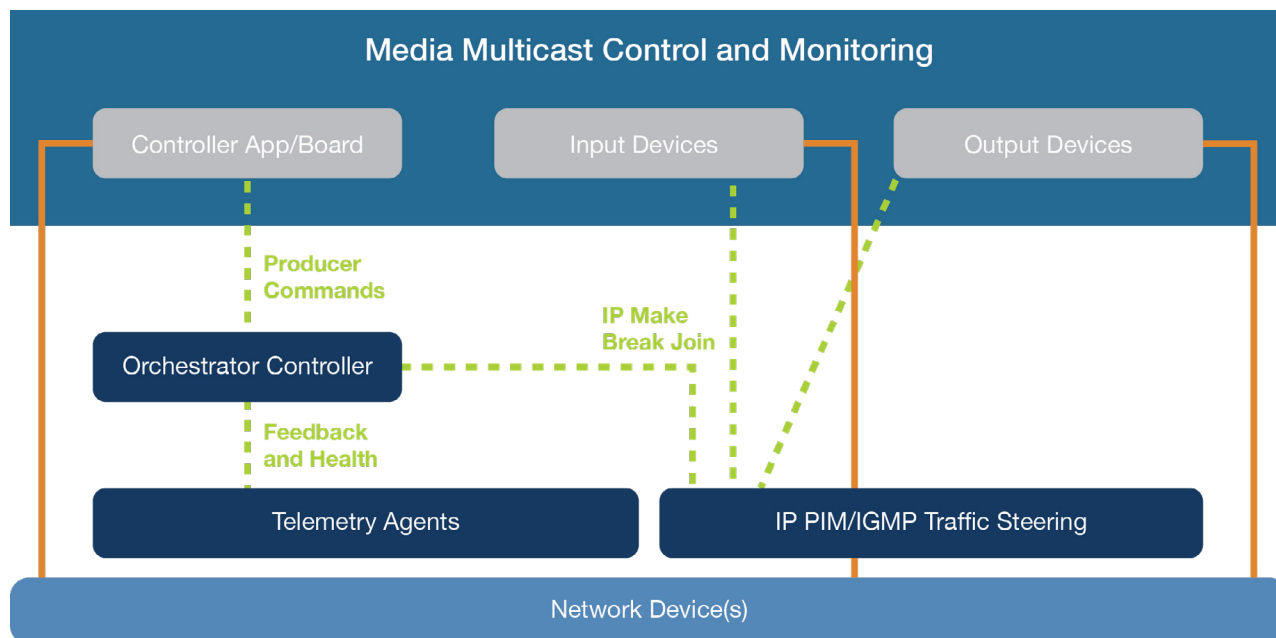
Control and Feedback

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.

While file based workflows use point to point TCP based file sharing protocols, broadcast workloads also require network based real time replication services, like multicast, plus traffic steering APIs so 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, recorders and others. Fast and reliable IP multicast joins and leaves replicates the behaviors of SDI routers. Inconsistent or slow multicast behavior compromises application performance and the broadcast workflow risking the end product.

Another key to optimizing the user experience are application programming interfaces (APIS) that help connect user actions on broadcasting equipment to network traffic actions on the underlay IP network. SDN orchestrators are the middleware that relays actions executed on a video producers control board to the network operations that fulfill those actions. Video source joins, merges, patches, fades 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 control boards, cameras, playback devices, compositors and other devices in the video production network. The reliability and confidence of a broadcast studio depends on the quick and faithful execution of the networking gear, their APIs and their interaction with the SDNO. The producer expects a near instantaneous result from his control button push. Any less is unacceptable.



Standardizing APIs in an IP-based broadcast studio

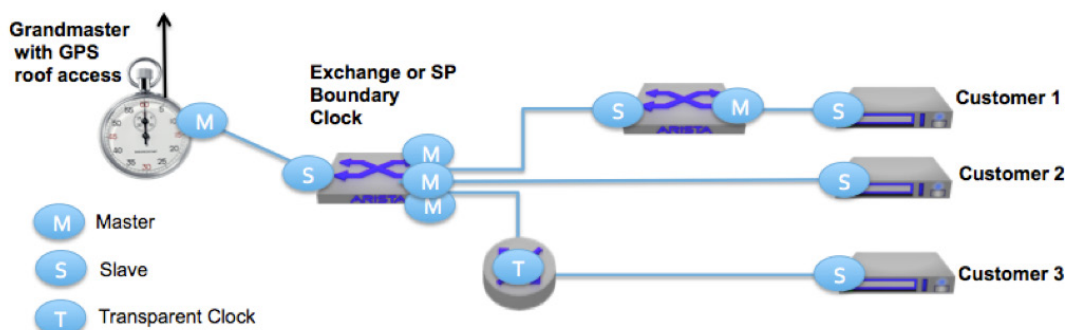
An uncThe variety of connected devices used in broadcast applications, coupled with the lack of standards for control APIs is troubling for consumers. Lack of interoperability or compatibility among proprietary implementations and future upgradeability can dampen the potential benefits of Ethernet based infrastructures. SMPTE and the VSF have worked hard to deliver standards and design frameworks for video and audio presentation, time synchronization and others. At the same time, the Alliance for IP Media Solutions (AIMS), a vendor supported interoperability group, has developed reference architectures for IP based broadcast networks to help clarify technical requirements for typical small and large broadcast studios.

While SMPTE and VSF have developed standards and reference architectures for IP based broadcast networks, network based orchestration is being addressed in another forum. The members of the Advanced Media Workflow Association (AMWA), a vendor supported alliance, have collaborated to define the framework for network controller APIs in its Networked Media Open Specification (NMOS) framework. Device discovery, registration, management controls and other key functions are now specified under their IS-04/5/and 6 API frameworks. These documents are helping vendors to normalize services with the objective of eventually ensuring not only compatibility among different platforms, but also open opportunities for interoperability and collaboration of devices.

Timing and Synchronization

Platforms used in broadcast 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/television monitor officials, help underscore the importance of timing synchronization. One simply can't lose a critical video frame in these critical situations. Similarly, the packetization of video frames over a 10Gbps or faster networks requires timing precision measured in microseconds with timing jitter measured in the tens of nanoseconds.

MPTE 2059-1/2 has adopted the IEEE-1588 precision time protocol as the service for time synchronization in IP based broadcast Ethernet networks. PTP is IP based and provides services for timing synchronization and drift correction among master, boundary, and transparent timekeepers. PTP sources, typically GPS controlled grand master clocks, provide timestamp packets to clients. Client devices reply back to account for accumulated time drift due to network path delays. In small studios, synchronization can be accommodated with a single source clock communicating to all the studio's time dependant devices. However in larger studio networks, PTP processing doesn't scale well to a single clock so it's good practice to distribute time synchronization duties among slave/boundary timekeepers.



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 grand masters and provide timing and drift information to upstream and downstream devices. In contrast, transparent mode PTP time keepers act as a helper, adding time drift information but not acting as a proxy to communicate directly to end devices. Distributed time synchronization with boundary mode offers some resiliency 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. Since transparent mode devices don't interact with end stations, they can't provide timing support in case of a master clock failure.

Telemetry

An uncomng network health and performance is important for any workload. 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 in speedy response to queries, and lacks flexibility in programming APIs. SNMP is built in a batch model: a MIB tree of data is delivered within seconds of an operator's query. It's not designed to provide split second updates. Maintaining SNMP based tools is also cumbersome because MIB schema must be known and compiled beforehand. There's no no built in key-value dictionary information that would help programs implement new or adapt to reordered data structures dynamically.

Developments in networking hardware, coupled with new programming design principles implementing model-view-controller programming paradigms, have greatly expanded the scope, power and responsiveness of network telemetry. The publish-subscribe model in the MVC paradigm lets networking and management devices reduce their workload by only exchanging data that has changed in state. So now, network status updates occur in hundreds of millisecond intervals while reducing load to network and monitoring devices. New programming paradigms also provide self defining key-value dictionary data structures. This data is presented in Javascript Object Notation (JSON) format: a common paradigm supported through virtually every modern programming language. The benefits to operators is that monitoring tools and SDN orchestrators can utilize near real time information streaming to provide enhanced monitoring, alarming and control feedback to users and programs. An example in broadcasting would be IP based broadcast controllers monitoring multicast throughput in near real time, to automatically control video streams through tie-line interfaces. With telemetry improvements the network working with SDN orchestrators automatically control broadcast workloads so producers can focus on the broadcast instead of the network.

Data center Architectures enable Animation and Post Production:

This new world of digital production has forever changed the CGI landscape. To produce realistic-looking, high-resolution visual effects (VFX) and animation, modern CGI shops are investing in cloud-scale parallel computing, with hundreds of servers driving tens of thousands of compute cores. These HPC clusters include centralized storage systems and file caching accelerators, all interconnected through high-speed, ultra-reliable switching.

The increasing dependence on the digital production workloads is raising the requirements for better reliability and non-stop administration and upkeep.

Aggregating bandwidth requirements for I/O heavy file-oriented post-production workflows such as feature-length films or episodic TV, or the low latency requirements of real-time workstreams like live, or near real-time edited broadcasts, generates a stringent requirements list that can be fulfilled by premiere networking providers. This enhanced paradigm creates new opportunities for innovative, specialized teams to compete in an industry that was unreachable except for the most

heavily funded players. The studios that can innovate, using these new digital technologies to produce high quality content in the shortest time frames, are the ones who win and prosper.

While high-performance computing was previously affordable by only elite research and development communities, the cost of this technology has dropped significantly, to the point where high-performance parallel computing infrastructures are cost justified in many applications. These include: automotive safety testing, oil and gas exploration, drug exploration (bioscience research), financial trading and cloud hosting. Media and entertainment companies are also rapidly embracing these technologies.



Standardized Editing Stations



The similarities between Render Farms and HPC datacenters are striking

Digital Media Technology Shifts

Smaller digital studios and post-production facilities leveraging these technologies are competing aggressively with larger organizations. These smaller studios produce HD or higher quality movies, ads, and special effects in many different streaming and playback formats. Remarkably, they deliver these projects with smaller budgets and in less time. Studios are also outsourcing their workloads to the cloud. These production houses have refined their business models and workflows, understanding the costs/benefits of outsourcing compute in off-peak hours, economic costs for offsite CPU and storage, and the development and administration costs for managing cloud-bursting. These organizations not only leverage high-performance compute and outsourcing, but they've also discovered the cost savings of automated administration and management tools that ensure the flawless execution of rendering/CGI workloads.



Today, the playing field is leveling due to an evolution away from expensive, proprietary systems, to more open and commodity-based platforms. This migration delivers dramatic price-performance advantages along with remarkable ROIs of operating technology. In the networking technology area, for example, maintaining an isolated Fiber Channel network for storage is unnecessarily costly, when there's alternative converged Ethernet technology options delivering 10Gb, 25Gb, 40Gb, 50Gb and even 100Gb that is interoperable and available from a variety of vendors.

Innovative studios are benefiting from these technologies and are deploying cloud-grade networked compute clusters within their data centers. Below, we outline seven of the top technology requirements:

- Centralized storage repositories hosting media libraries for content repurposing and more structured management of their assets. The most common storage deployed is file based, using caching or parallelization technologies with multiple 10Gb or 40Gb Ethernet interfaces to support large infrastructure and network fan-in designs.
- Rendering applications for post-production, sophisticated 3D, visual effects (VFX), animation, and finishing. These rendering applications use interconnected servers with high-performance 10Gb or 25Gb Ethernet switching platforms, requiring higher bandwidth 40GbE, 50GbE or 100GbE aggregation interconnects within the data center. The switching infrastructure must comply with data center class requirements for power and cooling efficiency along with cabling and operational automation requirements.
- Fast movement of large volumes of data (master frames) between artist's workstations and centralized storage or compute. This includes options for very low latency, sub-microsecond, networks or buffered infrastructures with virtually lossless traffic characteristics. These options offer an optimized solution for workflows like remote rendering and Virtual Desktop PC over IP technologies (VDI/PCoIP), or near real-time editing of broadcast workstreams. Either use case requires high-speed Ethernet switching using cost-effective twisted pair cabling and deep packet buffers to handle network speed changes between office workstations and data center server/storage nodes. Typical wiring closet switches are not built for these types of file transfers.

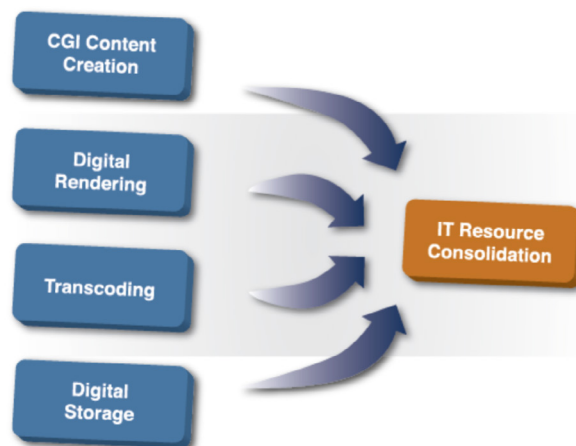
- Boundary Mode IEEE-1588 PTP timing services to enable alignment of real time broadcast work streams across the same infrastructure.
- As subscribers take advantage of content delivery across many device types, studios must deliver concurrent transcoding to multiple formats on the fly. This requires significant parallel server processing like what is available through data center clustering. In some scenarios, the data rate exceeds 10Gbps. Fortunately, with the recent availability of industry standard 25Gb Ethernet, there are more economical options to pre-existing layer two- or three-based load sharing schemes using 10GbE. Scalable networking options help ensure that networking infrastructure can keep pace with evolving formats and help deliver these derivatives in real-time.
- Near real-time telemetry, using advanced publisher/subscriber APIs so network control applications and operations personnel can identify and resolve issues before users do.
- 7x24 uptimes for all these use cases including storage, rendering, digital production and transcoding. CGI, rendering and post-production work are now real-time/all the time services. As studios work around the clock to get their final products delivered to many different types of global markets, their IT infrastructure must evolve to meet increasing demands and deliver constant up-time. As a result, studios are redesigning their infrastructure, moving many of their processes into modern data centers and using modern networking architecture. Administrators are leveraging high availability, active/active L2 features as well as L3 load sharing capabilities in the DC network. They are also expanding their management systems by using automation tools to support increasingly large and complex workloads while also minimizing the risks associated with human error. Reliability is greatly improved, while provisioning and reconfiguration are also expedited to improve productivity and reduce turnaround time.

High-Speed Networking for Computer Animation

Images created by movie animators, scene designers, and lighting artists are comprised of large data sets that must be composited and rendered to create a movie. To shave hundreds of hours of post-production time, the data is fed to hundreds of computers running in parallel; all interconnected with high-speed Ethernet through wire speed, non-blocking switches. Non-blocking Ethernet switching has become the de-facto technology of choice.

Similarly, images in modern animation are represented digitally via data sets. Like traditional animation, data sets representing animated cells are linked to comprise a full-length digitized movie. Animators, scene designers and lighting artists use computers to create and edit each data set that ultimately yields the visible frame. The pixel count of a 2K frame is approximately four times the size of an HD video image. The same frame rendered for 3D will have double the pixel density, or eight times the data of a HD frame! Modern digital movies are presented at thirty frames per second; faster than the frame rates of yesteryear. Uncompressed images can grow to dozens of megabytes in size and the actual movie grows to terabytes for the finished product!

The volume and size of animated frames and the increasingly complex rendering algorithms used to create more realistic images drives the need to deploy more compute servers for faster rendering turnaround. These increasingly high fidelity images also drive the need for higher capacity storage as 4K 3D cinema scales to 8K RealD 3DTM imagery. Hundreds of computers presenting tens of thousands of cores are orchestrated in a rendering pipeline to split the image processing process in parallel. This saves days, weeks, and often months of production time.




During production development, a full-length feature film with an average length of 90-120 minutes requires multiple petabytes of storage. For this reason, animated movie productions rely heavily on large numbers of multi-core servers, large file-based storage systems, and storage performance scaling with active data caching on flash and solid state disks (SSD). It only follows that the network also delivers wire speed high-performance with buffering to manage the flash congestion that occurs as hundreds of compute nodes aggregate their output to storage and caching systems.

The Need for Speed with Next-Generation X86 Based Workstations and Servers

The migration from proprietary CGI workstations to high-performance compute and rendering clusters is already underway. Virtual desktops leveraging PC over IP technologies reduces CAPEX costs per seat, protects intellectual property and offers new flexibility for moving and repurposing equipment and personnel in the studio. Networking infrastructure must be upgraded to accommodate these as well as other innovations, such as remote workstation rendering. Virtually all workstations now come with 1Gbps unshielded twisted pair (UTP) network connections, and servers come standard with 10Gb or 25Gb Ethernet connections. Given these performance options, network architects must accommodate this increased capacity with modern, multi-rate switches featuring ample data buffering. If neglected, new applications and platforms can easily overrun legacy network infrastructures, creating I/O bottlenecks that lose data and degrade the productivity of their creative team.

Leveraging advancements in chip technologies, Ethernet server and storage networking options continue evolving. Price/performance has dramatically improved with the advent of 10G Ethernet on UTP for compute and client, as well as the availability of 25Gb Ethernet server connections that is cost comparable to 10GbE Small Form factor Pluggable (SFP+) networking. New generation storage platforms offer 40Gb and 50Gb Ethernet services to support unprecedented IOPS, greatly improving the price performance of storage services. In summary, scaling performance in storage and compute, combined with distributed rendering and remote-authoring solutions is driving networking capacity requirements. Therefore, switching platforms offering a range of speeds including 1/10/25GbE and 40GbE connections with 100GbE interconnects is now a hard requirement. These offerings have reached industry maturity and deliver backwards compatibility with much better price/performance.



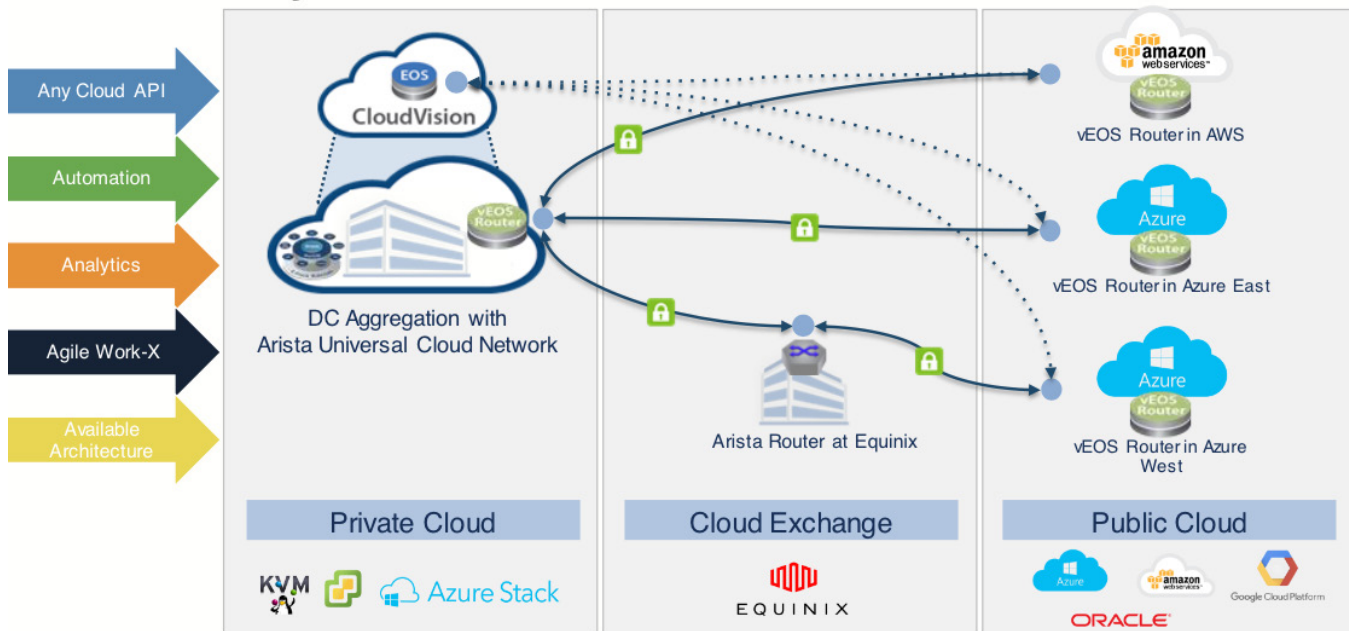
Arista Networks: Production Studios Preferred Choice

Arista Networks is the leader in the cloud data center networking market, with hundreds 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. Common amongst these applications is the requirement to expedite time to market for finished product. This means getting the work done faster, more efficiently, with less administrative overhead, and with no downtime or outage, all while being able to scale as more demands are placed on the network.

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 and 100Gb Ethernet switches defines the new standard for data center networking solutions. Arista's family of fixed and modular chassis 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. No other Network Operating System (NOS) has the adaptability and vendor support to make this feasible for users.

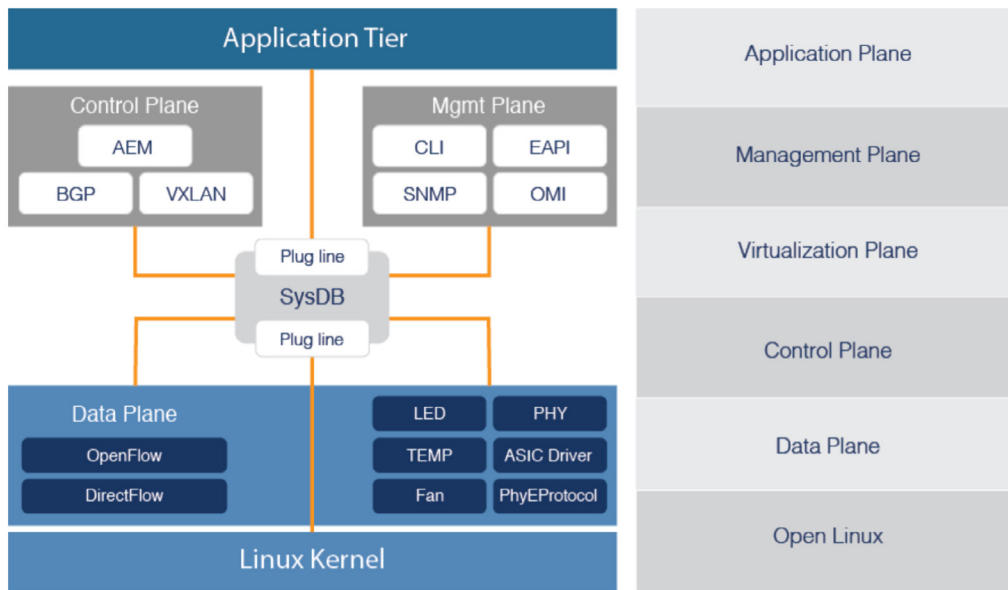
The power and modularity of EOS also lends itself to serving public clouds in the role of an aggregation virtual router. As cost and time to market needs dictate, content owners will deploy virtualized rendering or post production workloads across a variety of public clouds. The EOS vRouter provides network administrators a standardized routing virtual machine that's identical to, and can be managed like, all the other Arista data center 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, or KVM environment. Customers can model network architectures virtually using a version of EOS that runs as a VM. This is just one example of how Arista redefines the process and tools used to design, deploy and maintain modern media production data centers.

Arista Any Cloud Platform



EOS: Extensibility, Automation and Reliability

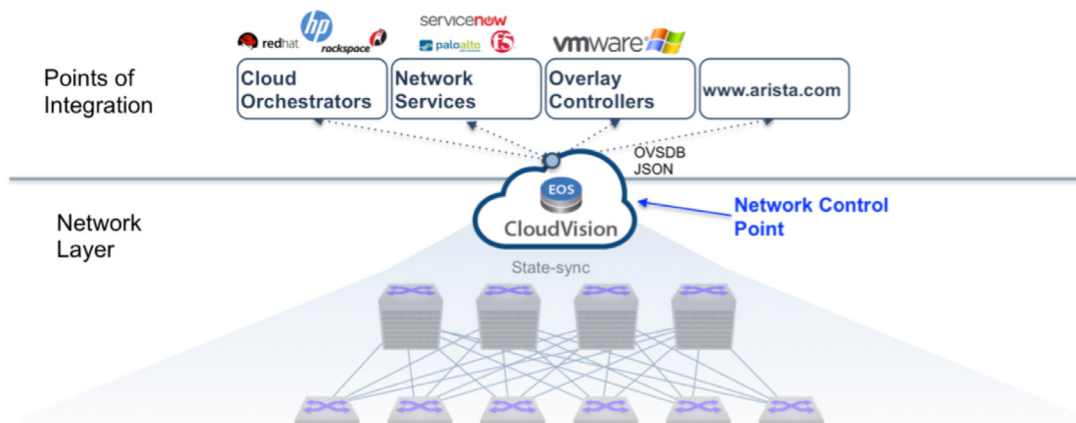
Arista's EOS is a fundamentally new architecture for cloud-grade media and entertainment workloads. Its foundation is a unique multi-process state sharing architecture that separates state information, protocol processing, application logic and packet forwarding. EOS system state and data is stored and 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.



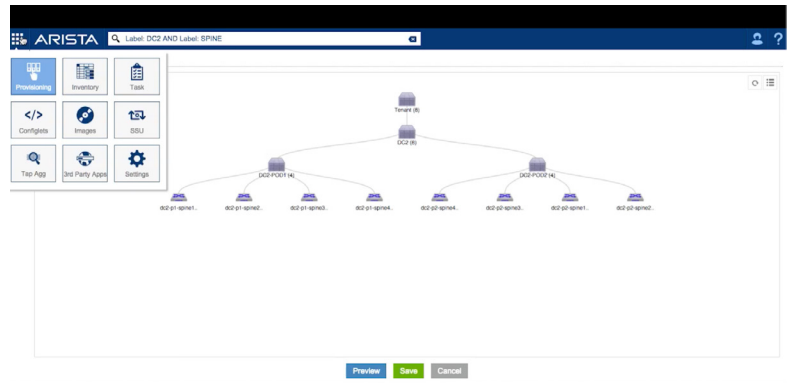
A cornerstone feature set, developed for IP based broadcast orchestration platforms, is EOS Multicast Control Services (MCS). MCS provides APIs that allow orchestrators to not only manipulate broadcast streams but also monitor them in near real time. MCS taps into multicast telemetry data that can be exported to SDN Orchestrators which in turn can display the telemetry or use it to automatically manage bandwidth for operators. MCS follows AMWA's API frameworks, allowing applications to ingest new or extended controls and telemetry using JSON data dictionary structures. MCS complements the existing open APIs in EOS that allows dozens of broadcast systems solutions providers to leverage Arista as a reliable transport for their networked broadcast workloads.

For automated change management, EOS has been built from day one 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 image remediation without service impact by using Smart System Upgrade (SSU) and ECMP maintenance mechanisms for faultless, and impact free management. Cloudvision also provides near real time telemetry monitoring, including multicast monitoring for broadcast oriented workloads.



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.

The diagram illustrates the Arista 7000 Switching Family architecture. It is divided into two main categories: 'Modular Spine/Spline™' and 'Fixed Leaf'. A central 'EOS' logo is positioned between the two categories. The 'Modular Spine/Spline™' category includes the '7300X Series' and the '7500R Series'. The 'Fixed Leaf' category includes the '7050X/7050X3/7060X/7060X3 Series' and the '7150S/7160/7280R Series'. The '7300X Series' and '7050X/7050X3/7060X/7060X3 Series' are grouped under 'General Purpose', while the '7500R Series' and '7150S/7160/7280R Series' are grouped under 'Value'. Images of the respective switch models are shown for each series.

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 and 100Gb for high bandwidth applications such as storage and inter switch connectivity.

All Arista platforms are purpose built with reliability in mind. All platforms have N+1 redundant field replaceable units, and are designed to run for many years. Last but not least, all Arista platforms run the same EOS binary image, which helps streamline qualification, reliability and maintenance. EOS also contributes to overall system reliability and maintainability by offering cloud proven wire speed layer two Multi-chassis LAG (MLAG) and layer 3 Equal Cost Multipath routing (ECMP) services. This ensures the data center infrastructure can increase throughput and reliability while faultlessly sustaining workloads even during the rare case of a link or system failure. Arista offers leading packet buffering capabilities to ensure traffic is not lost due to speed changes, congestion or microbursts. Studies demonstrate that buffering ensures traffic reliability, and optimized throughput, while providing an invaluable service for integrating legacy equipment to modern platforms without compromising the performance of either.

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 7500R 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 and 100G Ethernet connectivity. The 7500 implements a resilient Virtual Output Queuing (VOQ) architecture that ensures near 100% fabric efficiency without traffic loss 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. Finally the 7500R series modular chassis offers future proofing for customers who will ultimately require 400Gbps spine topologies.

Arista Solutions for Media and Entertainment

Modular 7500R platform for large-scale studios with an ability to deploy a 2-tier network architecture supporting scale-out broadcast and CGI/Post production:

- 4, 8, 12, and 16 slot platforms with bandwidths ranging from 38-150Tbps
- Virtual Output Queuing and deep buffers for uncongested low latency performance between compute and storage nodes
- N+1 redundancy for all system components
- High density 10G 25G, 40G, 50G and 100G wirespeed connections with future support for 400G
- 7200 series Deep buffer switches offering the same VOQ architecture as the 7500 but in economical fixed configurations ranging up to 60 wirespeed 100G interfaces
- Top of rack 7060X switch for high density 10/25G server and storage rack build outs
- Flexible platform portfolio for both low latency 10/25G twinax/fiber or 10GBASE-T structured copper cabling
- Non-blocking wire rate L2/L3 performance for meeting time-to-market pressures
- Low power, efficient cooling, and Zero Touch Provisioning for reduced operational costs
- Deep buffer top of rack 7020R or low latency 7010T platforms for connecting 1GbE workstations to large- scale workflows or real-time broadcast workstreams.
- Common Extensible Operating System (EOS) across all Arista platforms
- Arista Media Control Services (MCS) providing open, easy to use control and telemetry APIs for broadcast SDN Orchestrators.

- Software extensibility and APIs in EOS and CloudVision to deliver tailored monitoring, automated provisioning and change management
- CloudVision Multicast Monitor to provide near real time monitoring of network health and performance
- Smart System Upgrade (SSU) and MLAG features to support high availability and maintainability with sub- second infrastructure impact

Table 2: Workflow Types and Recommended Platforms

Workflow/Workstream	Requirement	Solution
CGI, VFX, Animation Authoring	1G UTP connection, handle large traffic flows, nearly lossless	7020R 1G switches with deep buffer VOQ architecture
Near real-time broadcast editing or compositing	1G UTP connection, low latency switching	7020R 1G switch with less than 3 μ sec end-end
Compute render platforms low bandwidth	Buffering for speed change and high speed uplinks	7020R 1G switches with deep buffer VOQ architecture
Compute render platforms high bandwidth	10/25G server connections and high speed upstream connections	7060SX 10/25GG switches with 40/100G uplinks
Highest performance render platforms	10G and 25G server connections with support for future proof scalability	7060CX and 7260CX Series 10/25/40/50/100G with low latency and wire speed performance
Content storage connections	High bandwidth and deep buffering for near lossless fidelity	Arista 7280R or 7500 deep buffer VOQ switches with 40/100G uplinks
Content transcoding systems	High bandwidth, low latency	Arista 7150S 10G/40G low latency switches. Arista 7060CX 10/25G with 40/100G uplinks low latency switches
Data center backbone connectivity	High bandwidth, virtually lossless, high density connectivity, 10-100G options	Arista 7500R modular platforms with deep buffers, VOQ architecture and connectivity scaling to over a thousand 10/25G ports and hundreds of 100G ports
Co-Lo Data Center Hosting	Scalable and flexible performance with a variety of interface speeds	Arista 7300X Spline systems for low latency and high performance
Virtualization scaling supporting VMware or OpenStack on the same DC infrastructure	Better, more economical scaling, sharing common infrastructure, reliable and self healing	Arista CloudVision VXLAN control services (VCS) supporting VMware, OpenStack and other OVSDB orchestration controllers
Infrastructure configuration management automation with self diagnostics and remediation	Automate change management, bug detection and remediation. Maintain code compliance without human intervention or production impact.	Arista CloudVision Portal. Provides tools to intake existing configurations, perform sanity checks, validate against master bug reports and recommend and implement remediation automatically.

Conclusion

The networking revolution in media broadcast, rendering and post-production CGI is opening new possibilities in digital realism, previously unimaginable imagery and cinematic storytelling. 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 the data center. Leveraging industry standard APIs, Dev-Ops tools or Arista's own media control Services and CloudVision helps expedite provisioning, changes and upkeep in the growing data center while containing administrative costs. These automation tools also avoid costly errors that may impact productivity. With Arista, media production companies can leverage 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



Copyright © 2018 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. Apr 3, 2018 09-0002-02