Powering Next Generation Video Delivery

Video Content Consumption Evolution

Consumer video consumption behavior has undergone a rapid transformation. The ubiquitous stationary TV has been replaced with far greater viewing options and an unprecedented flexibility to consume any content at any place on any screen. This revolution, first ushered in by the Digital Video Recorder (DVR) in the late 90s has culminated in a new world order of "binge" watching and freedom from the constraint of linear programming schedules or the anchor of the stationary living-room TV set.

Consumers continue to show a strong willingness to pay for flexible and unrestricted access to video content. Content providers, be it traditional video distributors (cable companies) or emerging on-demand video content providers (Netflix, Hulu, etc.), are aggressively monetizing this consumer trend. One important offering today is cloud-based DVR (cDVR) services. Traditional DVRs forced users to schedule recordings from home while limiting the number of simultaneous programs that can be recorded. By leveraging the cloud economics, flexibility, and scale, cloud DVR ushers in a new era of advanced video services allowing unlimited recordings, multiscreen video applications, and rich viewing experiences via 4K and Ultra High Definition TV.

Cloud Principles and Video Content Delivery

To meet the unprecedented demand for scale and performance, video content providers are evolving their infrastructure utilizing three key cloud networking building blocks: merchant-silicon hardware, open IP fabric, and software-driven provisioning and control.

- **Merchant silicon hardware:** According to estimates, by 2020, over 80 percent of the internet traffic will be video content, where the average video stream bit rate will increase three-fold to 16 Mbps. Video content will drive exponential network bandwidth growth. Merchant silicon platforms provide an effective CAPEX cost compression strategy to address the explosive growth of video traffic. To address this unprecedented growth, operators will have to leverage the benefits of Moore's Law, whereby network hardware processing power improves inversely to the cost of the network, thereby lowering CAPEX over time.
- **Open IP Fabric:** The Open IP Fabric delivers scale, performance and simplified network design through standards-based technologies. Open IP fabrics are based on two key components: two-tier, CLOS based leaf-spine topology, and an IP network stack based on open and proven protocols.
    - Leaf-Spine architecture: The leaf-spine topology has become the de facto standard for all-next generation scale-out networking topologies. It provides crucial support for a predictable non-blocking, any-to-any connectivity as well as scale and performance characteristics necessary to support a high performance video content delivery infrastructure.
    - Open and proven IP routing stack: By adopting standard IP routing protocols like BGP, open IP fabrics replace flat, complex, and non-scalable Layer-2 topologies with Layer-3 leaf-spine topologies. In addition to removing complexities that burden Layer-2 topologies, IP based fabrics eliminate expensive and complex vendor lock-ins and provide a resilient equal-cost multi-path routing (ECMP) topology.
- **Software-driven provisioning and control:** As part of the transformation to an IP-based video architecture, content providers are absorbing an ever-increasing number of network elements. A content video data center footprint can consist of thousands of compute, storage, and network devices. Automation simplifies and optimizes large-scale infrastructure management while avoiding costly human errors. Automation can be leveraged, at scale, for device deployments, monitoring, change management, troubleshooting, and software upgrades.

The Arista Advantage

Emerging video distribution services like cDVR have complex traffic profiles that can benefit from Arista's design framework and innovative platforms. Arista's network design is based on three transformational approaches to cost-effectively deliver a rich and reliable video experience:

- Technology transitions leveraging merchant silicon
- Resilient & high performance open IP fabric
- Operational efficiency via automation and enhanced visibility

Leverage Merchant Silicon

Arista's merchant silicon-based hardware designs are purpose built to help operators seamlessly address the increasing cost-to-scale challenge. The Arista 7500 Series of modular platforms has been tripling in density and performance every three years, with the 7500R3 series representing the fourth generation of the platform.

Figure 1: Industry leading 400G offerings with Arista R3-series

With over 460Tbps switching capacity and support for 288 400GbE ports the 7800R3 Series routing platform not only triples capacity and scale but also manages to dramatically improve functionality, increase IP forwarding table sizes, and enhance buffering. Leveraging Moore's Law, Arista's merchant silicon strategy will continue to deliver improved scale and bandwidth while reducing price per port. While the traditional legacy networking industry has had mixed results in delivering investment protection, the Arista 7500R3 and 7800R3 family delivers unprecedented multi-generational interoperability and investment protection across all components that make up the Arista R3 Modular Series.

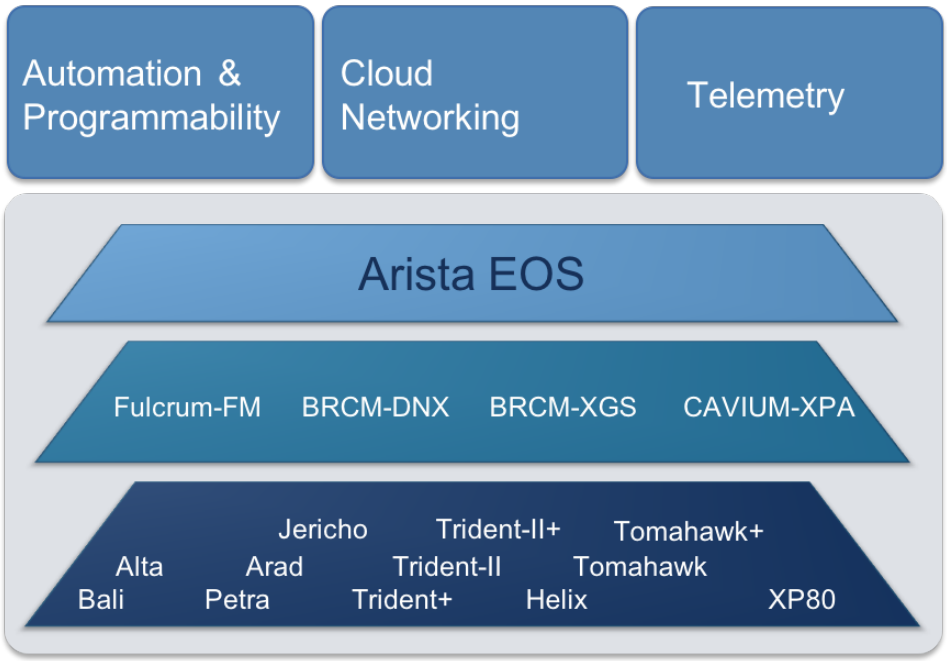

Figure 2: Arista EOS - Industry's most versatile open Network Operating System

Arista's Extensible Operating System (EOS) is designed from the ground up to uniquely support merchant silicon. In particular, EOS is the only networking software with support for multiple families of merchant silicon architectures, to optimize price performance and feature innovation. Arista has championed the use of merchant silicon and open-source software to sustain a significant pace of innovation by delivering a **single-binary-software image** that runs across all of Arista's products.

Resilient and High Performance Open IP Fabric

The quality and richness of the cDVR end-user experience is highly dependent on the performance and availability of the network. The cDVR network layer is required to support a set of unique traffic profiles. Unlike simple on-demand video streaming services, cDVR systems are susceptible to acute concurrency. Concurrency is the number of subscribers that can simultaneously access the service. During prime watching time cDVR platforms can experience up to 40 percent recording concurrency. Concurrency is further exacerbated by play-backs. Typical users replay recorded programs within the first 72 hours of the initial recording.

To meet these demanding use cases, Arista has championed non-blocking network performance and constant cross-sectional bandwidth availability through the use of open standards-based network fabrics with unmatched density and power efficiency. Building on merchant silicon innovation and standards-based network design, Arista advocates Layer 3 leaf-spine fabrics that are massively scalable to support line-rate multicast and unicast forwarding.

While IP Layer 3 ECMP fabrics have emerged as the status quo design choice, Arista's leaf-spine IP fabrics provide industry-leading features that offer a superior solution:

- Industry's leading symmetric unicast and multicast forwarding scale/performance

- Stateful and programmable infrastructure
- A deep buffer network fabric
- Superior availability and resiliency

Symmetric Unicast and Multicast Forwarding Performance

The performance of the cDVR service is highly dependent on consistent and predictable network performance. In particular, the delivery of a quality cDVR experience will require a high performance multicast fabric. Continued advancement in the delivery of a rich screen experience, with 4K & UHD streams, has pushed the pixel rates of video streams. Forecasts indicate emerging standards such as 4K video recordings and play-back will consume 2-3x the current bandwidth.

Operators deliver traditional linear video traffic (e.g cable TV programming) over multicast. The multicast forwarding is typically proportional to the available number channel lineups. Hence, every additional subscriber has a somewhat linear traffic impact on the video processing plants. On the other hand, cDVR relies on multicast traffic to acquire the user specified channels but utilizes unicast to playback the recordings to the end users. However, for each subscriber, the same recording can be played back to multiple devices (tablets,phones,TVs). Hence, cDVR unicast streams are loosely proportional to the number of end-user devices rather than the number of unique subscribers. This usage pattern can dramatically drive network bandwidth and performance requirements.

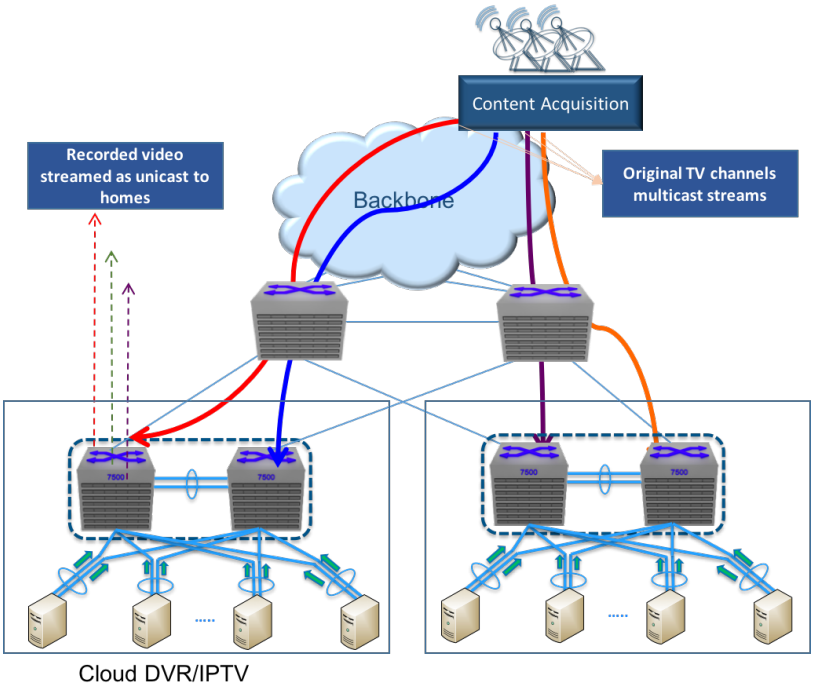

Figure 3: Next Generation Multicast Architecture

Arista network platforms seamlessly scale to support simultaneous line-rate forwarding of unicast and multicast traffic with low latency. Today, Arista platforms can support line-rate multicast and unicast performance at 10G, 25G, 40G, 50G, 100G and 400G interconnects with up to 460 Tbps of switching capacity. Arista based IP fabrics guarantee line rate throughput of multicast traffic, along with support for large multicast routes and fast IGMP join/leave rates.

High-Performance IP Fabric with Deep Buffers

Concurrent video recording and streaming on top of the hour, that correspond with program start times, are unique hallmarks of cDVR traffic patterns. High levels of concurrency can lead to sustained microbursts resulting in packet drops. Packet loss ultimately degrades the viewing experience in the form of video tiling and frame skips.

In the United States, regulatory requirements force operators to store a single unique copy per subscriber. For example, operators with 50,000 subscribers may be required to reliably store, retrieve and playback up to 15PB of video. To cost-effectively scale, cDVR deployments are migrating to a distributed commodity IP storage framework where data is striped across multiple storage targets. A common challenge in distributed storage network traffic is TCP incast, a many-to-one traffic pattern that can severely affect application performance. cDVR end-user playback requests can result in a severe incast condition where many hosts try to access the same storage device, or an intermediate host attempts to assemble striped video content from multiple storage targets.

Arista Network 7800R3, 7500R3 and 7280R3 series of routing and switching platforms are designed with unparalleled deep buffers to deliver an unprecedented 40ms of buffer capacity in both fixed and modular form factors. An IP/Ethernet storage architecture utilizing Arista's high-performance R3-series platforms combines deep buffers with an advanced forwarding architecture based on Virtual Output Queue (VoQ) to guard against packet loss from microbursts and incasts.

As with all high-performance networks, cDVR deployments require high speed interconnects, starting from the server and storage layer. The evolution of Ethernet access is undergoing yet another transformation with the availability of 25G/50G/100G standards. The Arista 7500R3 and 7280R3 series supports both Consortium and IEEE industry standard 25G, 50G and 100G. The 7280R3 can support up to 48 ports of 400G in a 2RU form factor, and the 7800R3 can support an industry leading 288 ports of 400G. This flexibility enables operators to build the high performance video delivery infrastructure, with Universal leaf and spine, providing access at all speeds, from 10G, 25G, 40G, 50G and 100G to 400G.

Highly Available Open IP Fabric

Quality and availability of video services are two key metrics that will drive adoption of popular and emerging content services like cDVR. Arista's resiliency solutions are designed for the 'always-on' infrastructure that remains responsive to the changing needs of operators. These solutions eliminate planned maintenance windows for software upgrades, minimize unplanned downtime and accelerate service velocity of cloud services like cDVR.

Arista IP fabrics achieve superior node and link resiliency via key Arista EOS technologies: *Multi-chassis Link Aggregation (MLAG) and Equal Cost Multipath Routing (ECMP).*

- Superior ECMP support with industry's fastest convergence times and the largest fan-out scale supporting up to 256-way active/active spine switches.
- Multi-chassis Link Aggregation (MLAG) multipath technologies provide active/active host connectivity to the leaf layer. These designs localize the L2/L3 gateway to the first hop switch allowing for the most flexibility in allowing different classes of switches to be utilized to their maximum capability without any dumbing down (lowest-common-denominator) between switches.

EOS builds on cloud networking principles of flexibility, simplicity and repeatability to deliver solutions crucial to support network resiliency and availability at the highest level. Arista EOS delivers high availability and resiliency at device, system and network level needed for service providers to deliver the rich video experience to consumers. Moreover, EOS software architecture delivers breakthrough resiliency solutions that include:

- Software Fault Containment (SFC) and Stateful Fault Repair (SFR): These features in combination enable faults to be localized to a single process preventing a system shutdown or reload and automatic recovery of software processes.

- In-service software patching and upgrades: Ability to upgrade individual processes without impact to switch operation, including packet forwarding.

Operational Enhancement via Automation and Enhanced Visibility

While scale and performance are key objectives, a rich end user service experience will require operators to enhance their service velocity, agility and availability. Automation is the key to improved service velocity and availability. Through automation, operators can expedite service provisioning, simplify and optimize day-to-day management and operations. Arista EOS delivers automation at all levels, from provisioning and monitoring to troubleshooting.

Provisioning and Management

To meet demand, video content delivery infrastructure will require a large number of compute, storage and network elements. Operators are locating their services closer to their end-users, which introduces the challenge of deploying and managing geographically distant infrastructure footprints. The ability to provision an existing or new greenfield network quickly and programmatically, at scale, can make the difference between a successful time-to-market and a lost opportunity. Arista EOS Zero Touch Provisioning (ZTP) automates the configuration of a new or replacement switch without user intervention. An extension to ZTP, Zero Touch Replacement (ZTR), enables switches to be physically replaced, with the replacement switch picking up the same image and configuration as the switch it replaced.

Monitoring and Troubleshooting

In order to truly reap the benefits of the cloud based transformation, operators will have to evolve their existing infrastructure management model to provide fine-grained real-time visibility. Current monitoring models based on polling (SNMP) and simple script based "screen scraping" cannot scale to deliver this aspect. Latency Analyzer (LANZ) is a breakthrough technology from Arista that tracks sources of congestion and latency with real-time reporting in microseconds. For example, microbursts can manifest in applications like cDVR that rely on high-throughput distributed processing that includes video encoders, recorders, packagers and storage targets. LANZ monitors and exports in real time the buffer utilization level on an Arista platform. During a period of congestion, LANZ can provide a time-stamped visibility to the event. This unique event-driven model allows congestion detection of microbursts as short as a few 100ns per queue level. For operators, this real-time visibility can help drive user experience by quickly resolving or rerouting around hot spots.

Legacy network visibility approaches have been based on inefficient polling mechanisms, which completely miss many critical network events that take place within a sub-second. Inconsistent and inflexible MIBs gather only a limited amount of actual network state. This creates a gap for operators in monitoring and troubleshooting process. Arista's NetDB provides real-time streaming telemetry via state streaming every system state such as configuration, counters, errors, route tables, environmental parameters, and much more. With this modern streaming telemetry architecture, customers will have access to both real-time and historic telemetry views of the network in one place and at a level of granularity never before achievable. This in conjunction with hardware-based accelerated sFlow provides a comprehensive visibility and telemetry solution for customers.

Conclusion

Consumer video consumption behavior presents operators with new revenue opportunities. The scale, flexibility, and variability of the new consumption pattern will challenge the existing architectures and management framework. Arista delivers key innovations for carriers to deploy next-generation solutions for Cloud DVR leveraging merchant silicon performance and power efficiency, open-standards-based technology, and an advanced, automated management framework. These include:

- Resilient operating system
- Best in class IP unicast and multicast forwarding
- Deep buffer platforms with flexible interface options (10G/25G/40G/50G/100G/400G)
- Congestion monitoring with LANZState-of-the-art real-time state streaming with EOS NetDB
- State-of-the-art real-time state streaming with EOS NetDB

Arista's continued innovation in best-in-class software, high-performance hardware, and advanced network management will deliver the scale and performance required by operators to support next-generation disruptive video services.

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