

CHANGING DYNAMICS OF IP PEERING

Arista Solution Guide

Inside

The Rise of Content Delivery Networks

Arista 7800R3 Universal Spine Platforms

- Highest density 100G and 400G Ethernet with power efficiency
- 4- and 8-slot chassis options

Arista 7500R3 Universal Spine Platforms

- High density 100G and 400G Ethernet with power efficiency
- 4-, 8-, and 12-slot chassis options

Arista 7280R3 Universal Leaf Platforms

- High density 10/25/40/50/100/400G Ethernet with power efficiency
- Broad choice of fixed systems

Arista R3 Series

- Internet scale routing
- Best in class convergence
- Deep buffer VoQ Architecture
- Programmable Traffic Steering with Arista EOS

Summary

At a time when the internet peering landscape has undergone a major transformation, the Arista 7800R3, 7500R3, and 7280R3 Series Universal Leaf and Spine platforms address the challenges of new IP peering dynamics with innovations in both hardware and software, disrupting the status quo economics for the content-driven internet.

Introduction

Looking back over the last two decades at the most impactful transformations in routing architectures, the internet itself has undergone one of the biggest changes. The primary driver for this has been the ever-accelerating growth of real-time content on the internet. From the 1990s to the early 2000s, internet content was primarily email, text, and static web pages. The network interconnects were symmetric or strictly hierarchical. Large Tier 1 Service Providers provided transit services to everyone else through their large global backbones, and smaller networks peered with each other due to the symmetric traffic flow. Fast forward to today, when content from various video streaming and social media platforms dominates internet traffic. This increase in high-bandwidth content has set in motion a series of changes to transit requirements, altering network economics, interconnection, and business relationships between the content owners, the Internet Service Provider (ISP), and the end consumers. New services are driving increased bandwidth, connectivity, uptime and latency demands, forcing network infrastructures to adapt.

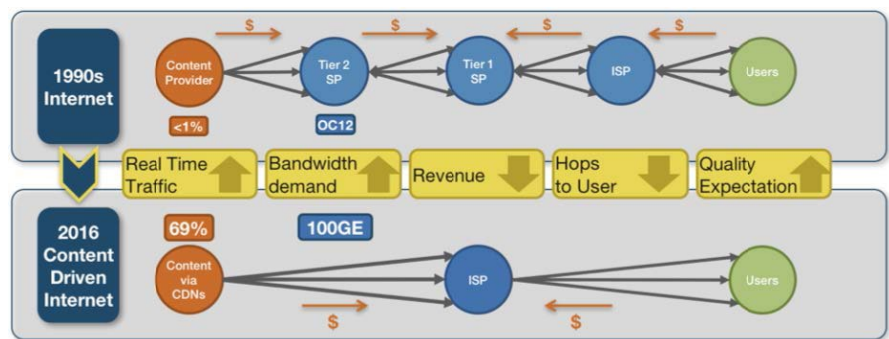


Figure 1: Evolution to Content Driven Internet

The Rise of Content Delivery Networks

According to the Global Internet Phenomena Report in October 2018 from Sandvine, almost 59% of downstream traffic on the internet is video and audio. In the world of content-driven internet, consumers are demanding a better entertainment experience for content that rides over the top of the ISP networks providing the last mile access.

In the race to capture eyeballs and provide a high-quality experience, most content providers are aggressively rolling out Content Delivery Networks (CDNs) that place content closer to the edge of the network to avoid transit hops through the internet and are peering directly over high-bandwidth peering interconnects with the ISPs. Content providers have several choices of exit points from their CDN (settlement-free public peering, direct/paid peering, transit, private WAN) and often rely on sophisticated algorithms to make the optimal routing decision and to traffic-engineer their content to meet the peaks in demand. In addition to the need for a high-performance peering platform, CDNs are often built in point of presence (POP) locations that are often severely limited in the amount of available space and power. This, in turn, highlights the value of both low power and compact footprint solutions for IP peering. Finally, Ethernet is the medium of choice for these interconnects, fueling the need for high density 10/100G, and in the future, 400G ports on the peering platform.

Meanwhile, the access ISPs are faced with increasing bandwidth demands on their last mile networks combined with declining per user revenues from their traditional fixed and mobile customer base. They face a set of growing competitive threats from new market entrants in the last mile access business that is further driving down end user prices. Also the largest CDN operators have offerings to house their servers in the data centers of the ISPs to distribute content directly from these local caches. All these transformations are forcing ISPs to refresh their network infrastructure to meet the growing bandwidth needs and adopt the cloud efficiencies and principles that have driven the evolution of compute, switching and storage.

The Arista Advantage

Content providers and ISPs have historically resorted to traditional router platforms to meet their internet routing needs. These platforms typically have a multitude of features and capabilities that are not required for an IP peering role, and are expensive, limited in port density, and very power hungry. Cloud providers have revolutionized the way the datacenter networks are built and operated using cloud principles such as scale out, simplify, and software-driven control.

The industry has now reached an inflection point where these established cloud principles are intersecting with the expanding capabilities of best of breed merchant silicon features and scale, and creating disruption in legacy routing system architectures. Leveraging merchant silicon capabilities with Arista EOS innovations enables a new best of breed routing platform. The Arista 7800R3 Universal Spine, 7500R3 Universal Spine, and 7280R3 Universal Leaf platforms offer content providers and ISPs a high-bandwidth, high density, programmable, power- and cost-effective option for both peering and cloud network infrastructure scale out.

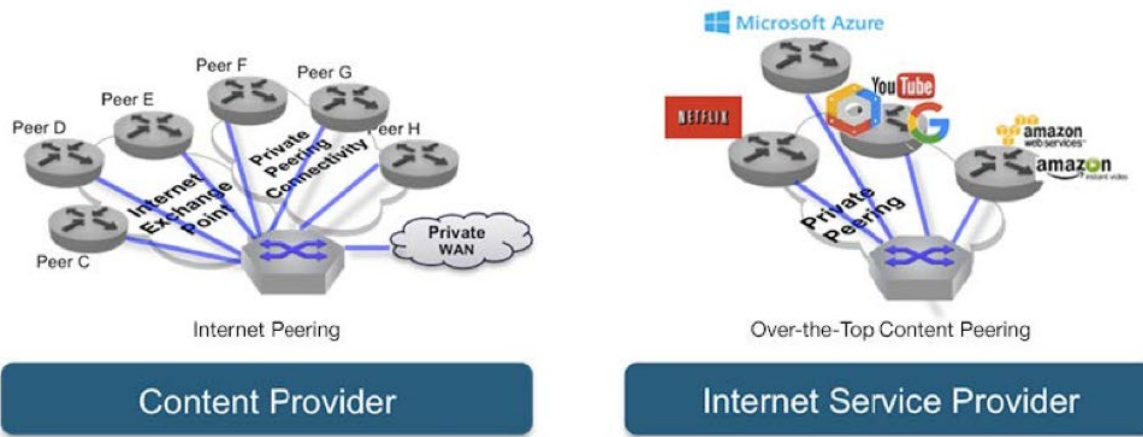


Figure 2: IP Peering Platform for Content Provider and ISP

Arista 7800R3 Universal Spine Platforms

Designed for the next generation of large scale virtualized and cloud networks the Arista 7800R3 Series modular switches are the 3rd generation of the R-Series universal spine switches. They combine 400G and 100G density with internet scale table sizes and comprehensive L2 and L3 features. Advanced EOS features for network monitoring, precision timing, and network overlays are combined with extensive hardware resources with a wide choice of systems and line cards to support lasting investment protection.

The Arista 7800R3 Series of modular switches deliver the industry’s highest performance with 230Tbps of system throughput to meet the needs of the largest scale data centers. With a choice of systems the 7800R3 Series provides capacity for four or eight line card modules, two supervisor modules, and six fabric modules. Designed for continuous operations and with system-wide monitoring of hardware and software components, simple serviceability and provisioning to prevent single points of failure the hardware architecture supports high-availability with hot-swap of all components.



Figure 3: Arista 7800R3 Universal Spine

The Arista 7800R3 Series leverages a deep buffer virtual output queue (VOQ) architecture that is 100% efficient, eliminates head-of-line (HOL) blocking and virtually eliminates packet drops even in the most congested network scenarios. An advanced traffic scheduler fairly allocates bandwidth between all virtual output queues while accurately following queue disciplines including weighted fair queuing, fixed priority, or hybrid schemes. As a result, the Arista 7800R3 Series can handle the most demanding internet peering and data center requirements with ease.

Arista 7500R3 Universal Spine Platforms

Designed for the next generation of large scale virtualized and cloud networks the Arista 7500R3 Series modular switches are the 3rd generation of the 7500R Series universal spine switches. They combine 400G and 100G density with internet scale table sizes and comprehensive L2 and L3 features. Advanced EOS features for network monitoring, precision timing, and network overlays are combined with extensive hardware resources with a wide choice of systems and line cards to support lasting investment protection.



Figure 4: Arista 7500R3 Universal Spine

The Arista 7500R3 Series of modular switches deliver the industry's highest performance with 230Tbps of system throughput to meet the needs of the largest scale data centers. With a choice of systems the 7500R Series provides capacity for four, eight or twelve line card modules, two supervisor modules, and six fabric modules. Designed for continuous operations and with system-wide monitoring of hardware and software components, simple serviceability and provisioning to prevent single points of failure the hardware architecture supports high-availability with hot-swap of all components.

Deep packet buffers and large routing tables provides complete deployment flexibility and allow the 7500R to be deployed in a wide range of open networking solutions including large scale internet peering applications.

Arista 7280R3 Universal Leaf Platforms



Figure 5: Arista 7280R3 Universal Leaf

Designed for next generation Cloud, Content Delivery, IP storage, leaf and spine networks and Data Center Interconnect the Arista 7280R3 Series of fixed switches combine a proven architecture with dynamic and deep buffering for lossless forwarding with high density 100G and 400G Ethernet, internet scale table sizes and rich L2 and L3 features that are space and power efficient for modern networks.

Wire speed L2 and L3 forwarding are combined with advanced features for network virtualization, open monitoring and network analysis, resiliency and architectural flexibility with deterministic performance. Deep packet buffers and large routing tables allow for Internet peering applications and the broad range of interfaces and density choice provides deployment flexibility.

Power Efficient 100G and 400G Ethernet

The R3-Series introduces support for high density 100G and 400G Ethernet in a variety of platforms from 24 ports of 400G in 1U to a 8-slot modular system with 288 ports of 400G, and increase performance by 3X over the previous generation.

A single port of 400G can support 1x400G, 2x200G or 4x100G modes with up to 36 ports in a single line card, equivalent to 144 ports of 100G. While demand for 400G is driven by the largest scale networks, the R3-Series enables high-density 100G deployments at lower power per gigabit.

The 7800R3 3kW power supplies are dual-input, leveraging an internal auto-transfer switch (ATS) to provide resilience in grid-redundant applications. This allows for an N+2 power supply deployment versus N+N, resulting in an increase in resilience, not power.

Consistent VoQ non-blocking fair delivery

The Arista R3-Series leverages a deep buffer virtual output queue (VOQ) architecture with a cell-based fabric that is 100% efficient and provides consistently low latency and jitter, eliminates head-of-line (HOL) blocking and virtually eliminates packet drops even in the most congested network scenarios. An advanced traffic scheduler fairly allocates bandwidth between all virtual output queues while accurately following queue disciplines including weighted fair queuing, fixed priority, or hybrid schemes. As a result, the Arista 7280R3, 7500R3, and 7800R3 series can handle the most demanding data center requirements with ease, including mixed traffic loads of real-time, multicast, and storage traffic while still delivering low latency.

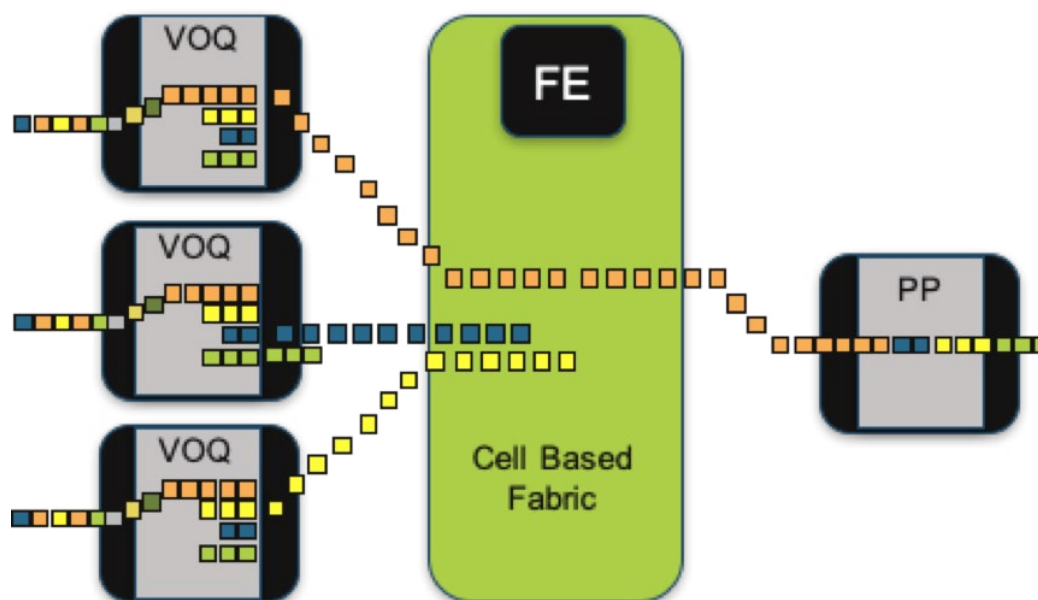


Figure 6: VOQ Architecture with Cell Based Fabric

The advantages of the VoQ architecture are that traffic is not forwarded to a destination without a credit, and that the cell-based fabric uses all paths to the egress packet processor to avoid creating hotspots. The cell-based fabric leverages a variable cell size and handles mixed interface speeds with no impact on system performance.

Enhancements for the R3-Series include multicast expansion in the ingress, fabric or egress, to reduce drops or queuing. In the event of a drop, it can be mirrored to a destination for analysis.

Arista FlexRoute™ Engine

All R3-Series products support full internet route scale with a capacity of 1.3M routes leveraging a common set of forwarding table resources for the L2 and L3 tables. The R3K-Series with expanded size tables can support up to 2.5M IPv4 and IPv6 routes, allowing for multiple overlapping full tables or for many years of growth of the Internet. Routes are held on the packet processor, without the need for external TCAMs or additional memory, affecting performance, power efficiency or overall system density.



Figure 7: Arista FlexRoute™ Engine

In addition to large table support, FlexRoute enables very fast route programming and reprogramming (tens of thousands of prefixes per second), and does so in a manner that is non-disruptive to other prefixes while forwarding table updates are taking place.

Arista Algorithmic ACLs

All R3-series platforms provide support for algorithmic ACLs that increase both the scale and flexibility leveraging the programmable pipeline. With support for over 100K rules, the R3-Series platforms are suited for internet edge and cybersecurity roles with stateful ACLs and rich counters identifying and filtering with no loss in performance.

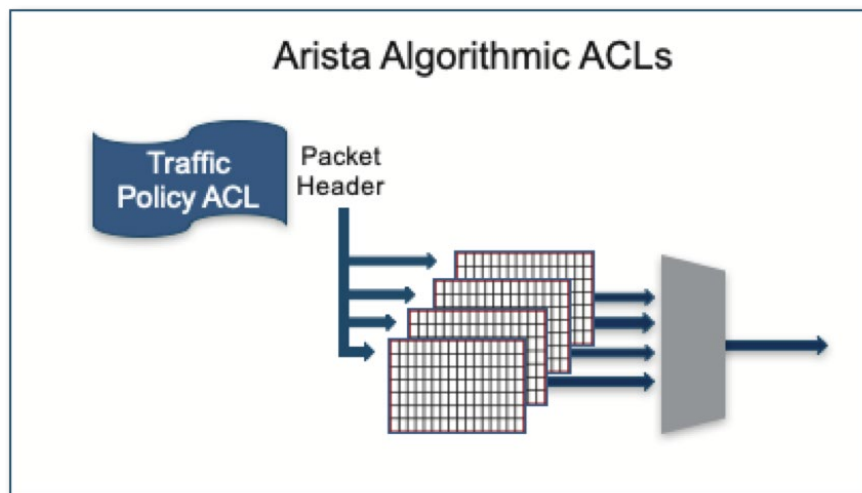


Figure 8: Arista Algorithmic ACLs

Algorithmic ACLs combine both software and hardware to enable more flexible and scalable solutions for access control, policy-based forwarding and network telemetry. Combining general purpose memory with advanced software algorithms delivers higher scale, performance, and efficiency with lower power and is more cost effective than traditional solutions. Algorithmic ACLs leverage efficient packet matching algorithms that in turn enables flow matching for access control, policy and visibility. The net benefits are a high performance policy engine with both increased functionality and scale in a cost and power efficient solution.

- Enables IPv4 and IPv6 access control at the same scale
- L4 rule ranges are programmed efficiently without expansion or reduced capacity
- Multiple actions can be performed on a single packet or flow
- User defined filters allow flexible packet classification based on offsets for custom actions
- Supports rich policy with consistent semantics that would exhaust classical resources

Flow Visibility

The programmable packet processing pipeline on the R3-Series platforms enables a range of new telemetry capabilities for network operators. In addition to new counter capabilities, flow instrumentation capabilities are enhanced through the availability of hardware-accelerated sFlow. As network operators deploy various tunnel overlay technologies in their network, sFlow provides an encapsulation-independent means of getting visibility into high-volume traffic flows and enables operators to more effectively manage and steer traffic to maximize utilization. The programmable pipeline provides these capabilities inline without requiring an additional coprocessor. Sampling granularity of 1:100 on 100G and 400G interfaces can be realized on all interfaces.

Programmable Traffic Steering with Arista EOS

All Arista products including the R3 Series runs the same Arista EOS software, simplifying network administration with a single binary image across all switches. Arista EOS is a modular switch operating system with a unique state sharing architecture that cleanly separates switch state from protocol processing and application logic. Built on top of a standard Linux kernel, all EOS processes run in their own protected memory space and exchange state through an in-memory database. This multi-process state sharing architecture provides the foundation for in-service-software updates and self-healing resiliency together with stateful switchover without the loss of data plane forwarding.

Arista EOS offers a variety of programmable traffic steering options for content providers to manage traffic across all available exit points from the CDN to optimally deliver content to the consumers, including MPLS, GRE, VXLAN, PBR to enable software driven exit point selection.

Conclusion

The internet peering landscape has undergone a major transformation in the past decade, driving the growth of CDNs, new business opportunities, and the refresh of access ISP network infrastructures. The content-driven internet of today demands a high bandwidth, cost-effective, programmable 100G/400G peering platform with low power and space requirements. The Arista 7800R3, 7500R3, and 7280R3 Universal Spine and Leaf platforms address the challenges of the new IP peering dynamics with innovations in both hardware and software, disrupting the status quo economics for the content driven internet.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

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

Vancouver—R&D Office

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

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

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

Singapore—APAC Administrative Office

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

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. Oct 2019 09-0001-02