Routing Feature Sheet

Arista 7280/7500/7800 platform

Overview

The 7280/7500/7800 R-Series platforms are a key component of Arista's Cloud Grade routing portfolio, the platforms are purpose built 10G, 25G, 100G and 400G systems providing industry leading performance, scale and port-density. The feature richness and flexibility of the platforms allows them to be deployed in a wide range of open networking solutions including large scale layer 2 and layer 3 cloud designs, the provider edge where scaleable L2 and L3 VPN services are required, in a traffic engineered MPLS core where high density 400G/100G is required, the metro-aggregation for the backhaul of E-LINE services, and in the Telco data center for EVPN overlay solutions. While delivering scaled out VPN services, with traffic engineering, the platforms maintain the inherent benefits of Arista's EOS operating system for real-time telemetry and network monitoring with open-APIs for automation. The R-Series platforms deliver a solution for operators with the performance and scale to meet on-going bandwidth growth, while reducing Opex costs by providing network-wide visibility and simplified automation, network orchestration and service enablement .

The Arista fixed and modular 100G/400G R-Series platforms

Cloud Grade Routing

Network traffic patterns and the year-on-year bandwidth growth in today's media rich cloud era demands a new class of routing platform and a new approach to network software in order to provide cost-effective on-going elastic scale while ensuring continuous operational and automation improvements at scale. To address this new industry challenge Arista's Cloud Grade routing portfolio is built on three fundamental principles.

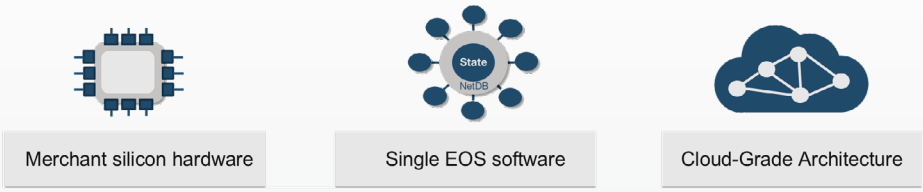

- **Merchant Silicon** - To provide a cost efficient and scalable routing platform, the proven approach by Cloud operators and now Service Providers is a scaled-out model with merchant silicon based hardware platforms. The market transition to merchant silicon based hardware is a fundamental principle of Arista's Cloud Grade routing portfolio with the R-Series platforms all built on merchant silicon. Over the past 12 years merchant silicon has successfully followed a Moore's law curve, where throughput and density has doubled every ~18 months, while reducing the overall cost and power consumption per bit to now successfully deliver cost effective 100G/400G hardware routing platforms. With consumers demanding ever richer media (4K to 8K) at an equivalent if not lower cost, a Cloud Grade routing platform built on merchant silicon is the only approach to provide a repeatable model that meets on-going bandwidth demands while ensuring a reduction in cost per bit.

- **Single Extensible Operating System (EOS)** - A key cornerstone to the success of next generation Cloud grade routing deployments will be the ease of automation, service orchestration and the ability to provide proactive cognitive closed loop network wide visibility. These are complex software challenges that are only amplified by traditional routing solutions which provide inconsistent APIs, OS architectures and telemetry models across a multitude of different hardware architectures. Arista's Cloud Grade Routing solution takes a unique approach to solving the problem by providing a single programmable state-drive operating system (EOS) that is supported across all Arista platforms including the R-Series routing platforms. This single EOS approach ensures a consistent automation, service orchestration and state-driven telemetry model, regardless of the platform or the platform's role within the overall topology. Thus as the network evolves and grows over time and next generation platforms are introduced into the solution, the automation, service orchestration and telemetry models remain consistent and new product onboarding is significantly simplified.

- **Cloud Grade Architecture** - To evolve from the traditional scale-up approach to routing design and the resultant costly rip and replace life-cycle, Arista's Cloud Grade Routing architecture is built on a repeatable modular scaled-out model that aims to reduce complexity in the protocol stack and hardware dependencies in the forwarding plane, while maintaining operational consistency via a single OS and API layer. The R-Series platforms support traditional VPN services (IP-VPN, PW etc) and transport models (RSVP-TE and LDP), but to reduce protocol complexity the platforms in combination with Arista's EOS software provide support for next-generation E-LINE, E-LAN and L3-VPN services using a single EVPN control-plane and SR-TE transport model. A major benefit of this control-plane simplification is the ability to now easily expand the forwarding plane based on bandwidth and service demands using a modular scaled-out approach, where disaggregated PE and P nodes are constructed using relevant sized merchant silicon platforms and open standard protocols. This open modular scaled-out design strategy provides the ability to now leverage the appropriate merchant silicon platform at each point in the network based on specific service demands, while enabling incremental upgrading or re-deploying of hardware as requirements change over time, without the need to rip-replace the entire infrastructure.

R-Series Deployment Models

The performance and port-density of the R-Series, in combination with the rich routing feature set of the EOS operating system, allows the R-Series platforms to be deployed in a wide range of routing use cases such as Internet peering, metro aggregation for E-LINE services, provider edge for rich Layer 2 and 3 VPN services and next-generation software driven traffic engineering solutions for the provider core.

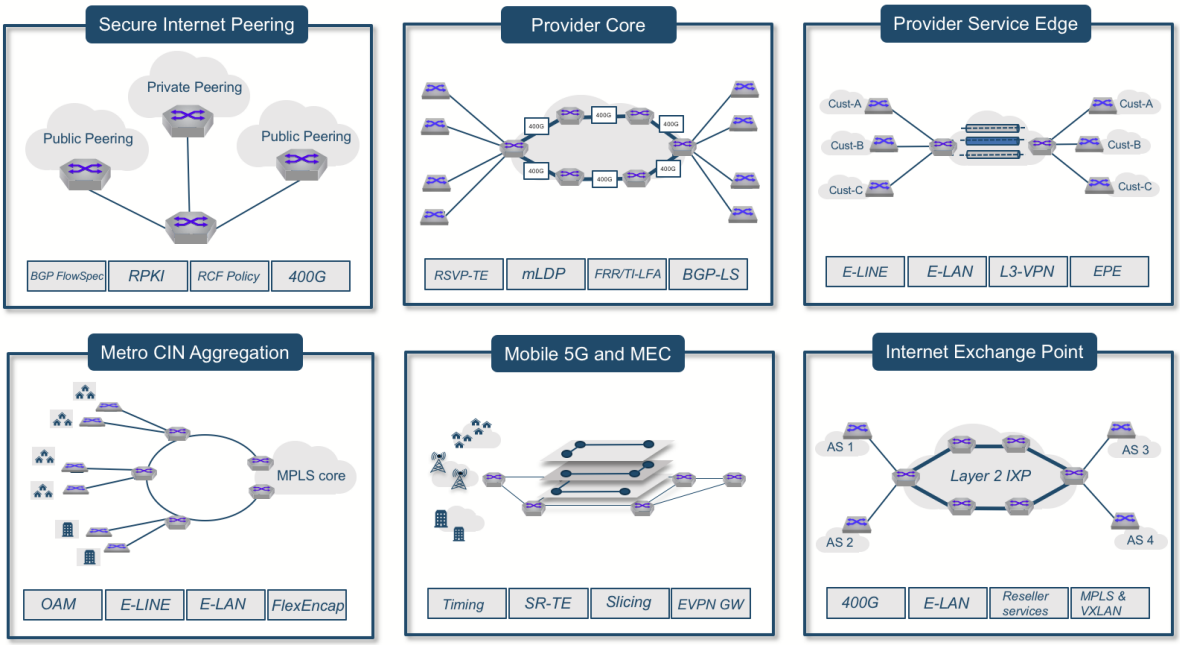

Secure Internet Routing and Peering

Continued increases in end user demand for Internet bandwidth drives a rapid and sustained requirement for capacity upgrades at Internet peering points. In addition, operators face a fierce demand to improve customer quality of experience. To support this twin demand for higher capacity and improved customer experience, operators are adopting a distributed peering approach that places content closer to the end users. The shift to a distributed model introduces a larger number of security threat points and DDoS targets. The Arista R-Series family of routers addresses these emerging requirements with a choice of highly secure 10G/25G/100G/400G compressed footprint platforms with a scalable RIB/FIB capacity that supports advanced peering features including:

- Arista's innovative FlexRoute Engine, which provides internet scale routing, allowing the R-Series platform to support 2.5M+ FIB routes and 30M+ RIB routes.
- Modern resilient EOS software control plane, which delivers high availability through in-service upgrades with best-in-class BGP scale and PIC convergence
- A rich security and DDoS mitigation feature set, including large modular ACLs with deep packet inspection, full featured Flowspec for DDoS protection, and a robust RPKI implementation for secure routing decisions
- Programmable and dynamic route policy control through EOS's unique RCF programming language

- Unprecedented fine-grained network visibility via a EOS's state-driven real-time telemetry, high capacity sFlow and hardware assisted IPFIX
- Advanced programmable management plane with a rich set of open API options: YANG models driven over Netconf, gRPC and REST interfaces

Service Provider Core

To construct the next-generation provider core, there is a requirement to aggregate higher volumes of 100G/10G uplinks from the provider edge, while providing high density 400G support for inter-core links. This needs to be achieved at scale to drive down capital cost, while maintaining support for traditional MPLS functions such as RSVP-TE, FRR, and mLDP as the core evolves over time to a state-free segment routing (SR) model. The Arista R-Series platform addresses these challenges at the provider core, by providing industry leading 100G/400G density in both modular and fixed form factors, while providing rich and scalable LSR capabilities for segment routing (SR), LDP, mLDP and RSVP-TE MPLS transports:

- The R-Series platforms are built on merchant silicon, thereby providing industry leading 100G/400G port density in the smallest silicon footprint, thus driving down power consumption while reducing the overall cost-per-bit.
- The rich EOS feature set of the R-Series platforms, provides support for multiple transport models (LDP, RSVP-TE, SR-TE) while providing a flexible label allocation framework to allow the coexistence and seamless migration of LDP/RSVP-TE deployments to next-generation SR/SR-TE solutions.
- Real-time network state telemetry and flow visibility through Arista EOS's state-driven architecture, to allow network wide visibility for operational troubleshooting.
- Open programmable EOS APIs (RESTCONF, NETCONF) for day zero deployment and on-going service enablement and troubleshooting allowing full life-cycle network automation and service orchestration.

Service Provider Edge

The provider edge is the critical component to delivering cost efficient revenue generating services, with operators challenged to deliver both traditional and next-generation services across varying port speeds (1G/10G/100G/400G), resiliently, at scale and in a power efficient form-factor. Equally, the speed of service enablement is vital to achieving a faster return on investment and real-time service visibility is essential to improving the overall customer experience. The Arista R-Series family of routers delivers the service flexibility and form factor required at the provider edge, while providing a state-driven EOS operating system to enable real-time service telemetry and programmable service enablement.

- The rich EOS feature set of the R-Series platforms provides full service flexibility with support for both traditional services (IP-VPN, E-LAN and PW) and next-generation EVPN services (L3VPN, L2VPN, E-tree).
- For mission critical services required to meet stringent SLAs, the R-Series support Fast ReRoute (FRR), RSVP-TE and SR-TE along with open programmable APIs (RESTCONF, NETCONF) for off-box tunnel provisioning and real-time state-streaming for intelligent path selection and topology discovery (BGP-LS).
- The R-Series, built on merchant silicon, delivers this rich service flexibility in industry leading 10G/100G/400G fixed and modular platforms, the higher port density driving down power consumption while reducing the overall cost-per-bit.
- Fine-grained network state and flow visibility via EOS's state-driven real-time telemetry architecture and high capacity sFlow and hardware assisted IPFIX implementation, enabling real-time network wide visibility for service visibility and troubleshooting.
- Open programmable EOS APIs (RESTCONF, NETCONF) for day zero deployment and on-going service enablement allowing full life-cycle network automation and service orchestration.

Metro and Cable CIN Transport Services

With the introduction of triple-play services, cloud computing, smart phones and IoT, the Metro network is under pressure from the on-going bandwidth demand and the expected always-on user experience. Further, to improve revenue and return on infrastructure investment, providers are no longer deploying independent mobile and FTTH access networks but moving to a new single hybrid CIN (Converged Interconnect Network) architecture. To keep pace with this change, the metro network needs to evolve to a programmable converged service model, delivering on-demand Ethernet and IP services at 10G and 25G speeds. To deliver these higher speed access services in a cost effective manner, there is also a need to aggregate a higher density of 10G/25G services in a smaller footprint for backhaul over a 100G/400G infrastructure. Arista's R-Series platform in combination with EOS's rich programmable Ethernet and IP services feature set, is purpose built to address the requirement of this next generation metro network.

- Resilient Ethernet services support, including E-LINE, E-LAN, E-Tree services (EVPN, PW, VPWS, L2VPN) in addition to layer 3 VPN services (IP-VPN and EVPN-MPLS)
- A Flexible encapsulation framework for addressing the complex packet manipulation requirements for delivering Ethernet and IP services. The framework provides support for QoS policies, layer 2 protocol transparency policies and VLAN tag manipulation actions (consume, push etc.) on a per service basis.
- A Rich OAM software stack for performance monitoring and end-to-end service and transport validation.
- To provide connectivity to remote and diverse locations. available space and power can be a major concern when deploying metro nodes, the R-Series built on merchant silicon ensures market leading port density, while delivering industry leading space and power efficiency.
- Real-time network state and flow visibility through the EOS state-driven architecture, with high capacity sFlow and hardware assisted IPFIX providing unprecedented fine-grain end-to-end network and service visibility,
- Open programmable EOS APIs (RESTCONF, NETCONF) for day zero deployment and on-going service enablement allowing full life-cycle network automation and service orchestration.

Mobile 5G backhaul and Mobile Edge Compute (MEC)

The roll-out of 5G services is expected to increase demand for data intensive applications. To drive revenue from the infrastructure, an operator's transport network needs to evolve from a pre-planned topology with static behavioural characteristics to an on-demand flexible transport layer that can be tuned and sliced to individual customer and application requirements (latency, bandwidth, resource etc), enabling the infrastructure to be truly application-aware. To meet these expectations, the transport network is evolving to a high-bandwidth 400G SR-TE architecture, where controller or controller-less approaches are adopted, based on cost and size, to enable dynamic network-slicing of the infrastructure based on application and service demands.

In this new application-aware infrastructure, to improve user experience by reducing latency, improving bandwidth and providing data offloading, there is a drive to push compute and storage resources closer to the network edge and therefore the consumer. The MEC (Mobile Edge Compute) model requires operators to build space-efficient, low-latency and high-bandwidth edge DC infrastructures, that can provide overlay (EVPN-VXLAN/MPLS), support for software based NFV applications, while providing seamless programmable connectivity to the SR-TE transport.

The flexibility and performance of Arista's R-Series platforms, in combination with EOS's rich programmable feature set, is unique in being able to address both the emerging demands for a flexible 5G transport and the building of mobile edge compute overlay environments on a single consistent platform, API and EOS operating system, with purpose built hardware and software features.

- To deliver the precise timing requirements for 5G service deliver the R-Series platform provides hardware support for fine grain timing through PTP (boundary and transparent clock) with PTP Telco profiles and SyncE*
- To deliver cost efficient and scalable network slicing solutions, the EOS software of the R-Series platform supports both a controller-less approach (BGP colored routes and FlexAlgo*) and controller based approaches with open APIs and state-streaming, for deterministic path selection and topology discovery (BGP-LS).
- A rich EVPN feature set for the constructing of low-latency MEC overlay fabrics, with seamless integration with the SR-TE transport via EVPN GW support and open programmable integration with third-party software overlay vendors (VMware/OpenStack)
- Open programmable EOS APIs (RESTCONF, NETCONF) for day zero deployment and on-going service enablement, allowing full life-cycle network automation and service orchestration across the infrastructure with real-time network telemetry for unprecedented fine-grain end-to-end network and service visibility,

Internet Exchange Point (IXPs)

Internet Exchange Points (IXP) provide resilient layer 2 infrastructures to enable BGP peering and the exchange of Internet traffic between Service Providers, Large Enterprises and Content Providers, thereby facilitating the sharing of services and content between public, private and enterprise networks. With content rich media driving bandwidth consumption, peering speeds at the exchange are transitioning from 10G to 100G and more recently 400G for high volume content providers, while at the same time IXP customers are demanding greater visibility, faster service enablement with ever more stringent SLAs. Arista's R-Series provides the scale and performance to meet these new 100G/400G bandwidth demands while providing a rich programmable EOS feature set to address the unique architectural and security challenges of IXP infrastructure.

- The rich software feature set of the R-Series platform, provides the ability to construct highly resilient, scalable L2 topologies using EVPN-VXLAN or EVPN-MPLS, with support for ARP suppression for BUM traffic reduction, and All-Active and Active-Standby multihoming for resilience,
- Highly scalable 512-way ECMP solutions enable straight forward bandwidth growth and path diversity within the layer 2 infrastructure, with support for resilient ECMP to minimise flow disruption during any link or node failure.
- EOS's Flexible Encapsulation (FlexEncap) framework provides support for switching, shaping and policing of sub-rate reseller services.
- Open programmable EOS APIs (RESTCONF, NETCONF) for on-going service enablement and troubleshooting, with full integration support within IXP Manager software suite.
- Real-time network state and flow visibility through the EOS state-driven architecture, with high capacity sFlow and hardware assisted IPFIX providing unprecedented fine-grain end-to-end network and service visibility,

Arista EOS state-driven operating system

As provider networks linearly scale and new revenue generating services are enabled, continuing to deliver an industry leading return on investment creates a need for a more programmable software-driven operational model that is both agile, while cost-efficient to ensure operating costs remain constant as the infrastructure grows. This evolution to a software driven operational model places greater focus than ever before on the architectural design, resilience and programmability of the software running within the infrastructure.

Arista's EOS is a next-generation state-driven modular operating system, designed to address the requirements of very large scale environments. The EOS architecture builds on a standard Linux kernel and runs all processes in their own protected memory space, cleanly separating switch state from protocol processing and application logic through an in memory database (NetDB).

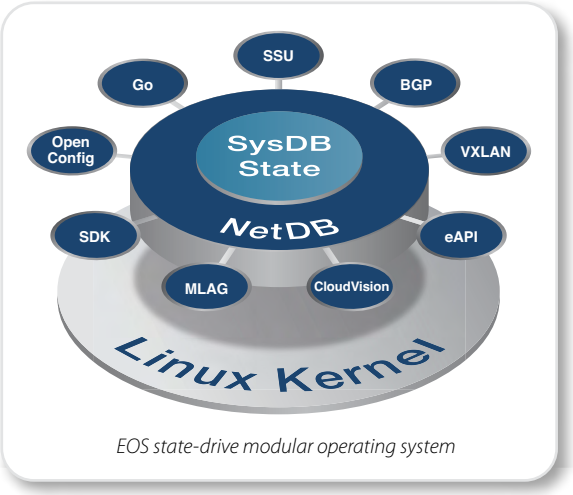

EOS state-drive modular operating system

The in-memory NetDB database (machine generated at run time) runs in user space and contains the complete real-time state of the system. Like traditional databases, NetDB does not contain any application logic and is only responsible for keeping state. However, rather than being optimized for transactions, NetDB is designed for synchronizing state among processes, also called 'EOS agents', by notifying interested agents when there is a state change. Each EOS agent subscribes to NetDB to be notified when the state of other related agents change within NetDB, When a state change occurs within an agent , updates are then published to NetDB, which then in turn notifies the subscribed agents interested in the change..

This centralized database approach to passing state throughout the system and the automated way the NetDB code is generated, reduces system overhead and simplifies inter-process communication to significantly reduce risk and error. By removing inter-process dependency and direct communication between agents the architecture also improves software feature velocity and quality , and provides openness for customers wishing to build their own applications, who can use the same in-built APIs to receive notifications from NetDB both for state visibility and feature customization.

Serviceability

The multi-process state sharing architecture of the EOS operating system, provides the foundation for industry leading availability and serviceability on the R-Series platform. In traditional network operating systems, software faults or security patches often require a software reload due to the monolithic architecture of the operating system, resulting in seconds to minutes of downtime. Reconvergence around such issues places additional load on neighboring devices, as topology changes ripple across the network.

With the modular architecture of the EOS operating system, a fault is contained within the agent or driver where the fault originated, in the unlikely event the fault causes the agent to crash, then the EOS process manager restarts a new instance of the agent. With the separation of state from the processing agents, there is no requirement for the re-starting agent to query or process any older stale state, instead the current state can be pulled directly from the NetDB database. If the fault causes the agent to hang or loop, EOS detects the condition and seamlessly restarts the agent, thus providing an architecture that becomes self-healing. This multi-process state-sharing architecture is also key to reducing maintenance windows by allowing more operational tasks to be performed during normal l hours, without the need downtime; EOS agents can be patched live and restarted if necessary without disrupting the overall operation of the node.

Programmability

As provider networks linearly scale and new revenue generating services are enabled, continuing to deliver an industry leading return on investment creates a need for a more programmable software-driven operational model that is both agile, while cost-efficient to ensure operating costs remain constant as the infrastructure grows. This evolution to a software driven operational model places greater focus than ever before on the architectural design, resilience and programmability of the software running within the infrastructure.

To provide a new software-driven operational model that is both agile, while cost-efficient EOS is fully accessable across all layers of it's software stack – Linux kernel, hardware forwarding tables, node configuration, control plane as well as the management layer via open APIs, the rich set of structured APIs includes:

- eAPI JSON based RPC, providing a REST-like interface for configuration and monitoring, using native CLI commands
- OpenConfig, Go, Python and Ruby based object models
- Native Go and Python on box scripting
- Develop native high performance applications using EOS SDK
- Native Linux APIs and scripting
- Tight dev-ops integrations with Puppet, Chef and Ansible
- NETCONF and Restconf transport protocol

The open programmability of the EOS software stack, and the storing of state in a common database, NetDB, which is easily accessible through the open APIs, allows the the R-Series platforms to be rapidly and easily integrated with a wide range of third-party and open-source applications for service initiation, traffic engineering, network management, automation, and network orchestration.

Real-time state streaming and analytics

Arista EOS software supports the traditional Syslog, SNMP traps and polling mechanisms for collecting and reporting routing state, however these traditional approaches can often be restrictive when trying to monitor protocol state, environmental alarms, memory, buffers etc. in real-time. To provide this level of visibility, Arista's EOS software supports real-time state telemetry. All network state (interface statistics, configuration, protocol, routing table, environmentals etc.) is stored within EOS's state-driven NetDB database, and can be streamed off-box in real-time via gRPC. This open standard interface provides third-party monitoring tools and Arista's CloudVIsion Portal (CVP) an unprecedented level of visibility into the ongoing health and performance of the routing infrastructure.

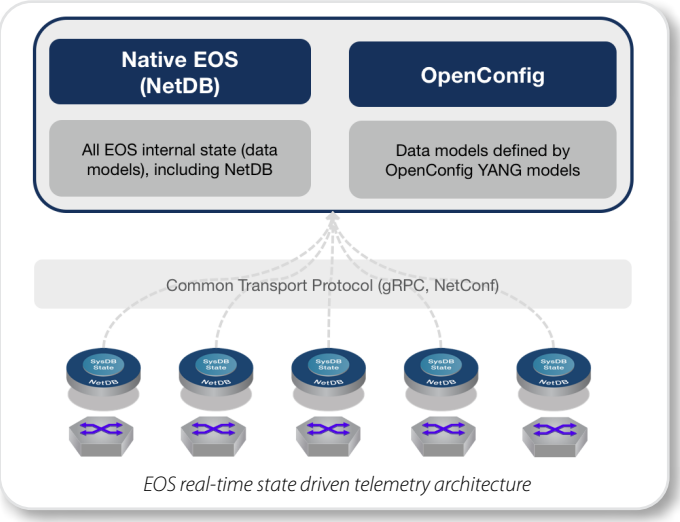

EOS real-time state driven telemetry architecture

With the state collection, through CloudVision or third-party tools, based on real-time streaming rather than polling, state updates are provided continually as changes occur. This makes the approach vastly superior to traditional polling models that have a typical granularity in the order of minutes and are often limited due to system control plane capacity. Also unlike traditional pre-defined SNMP MIB-based approaches, the EOS state streaming model is capable of streaming the full state of each node in the infrastructure, this includes details such as configuration, counters, errors, statistics, tables, environmentals, buffer utilization, flow data, and much more. This streaming architecture is also the basis for EOS's YANG-based OpenConfig data models.

Arista Routing Architecture

The modern state-driven architecture of Arista's EOS operating system, in combination with the performance and scale of the R-Series hardware platforms, enables a more programmatic Cloud based approach to building next generation routing architectures. This facilities more efficient use of hardware resources (ACLs, routing policies, FIB tables), faster service enablement and greater visibility (open APIs and real-time telemetry), all with within a repeatable modular scaled-out architecture that aims to reduce complexity while ensuring continuous operational and automation improvements at scale.

Routing Control Functions (RCF)

Deploying routing policies and route filtering at scale is a cumbersome and complex challenge due to the limited software constructs (permit, deny, continue, goto etc) available within traditional network operating systems. Arista's EOS Routing Control Function (RCF) addresses this complexity by delivering a powerful programmatic approach to large scale route policy definitions and attribute modification. This new programmatic approach means traditional complex multi-line, multi-peer routing policies can now be defined as discrete RCF functions, where an RCF function could be applied across multiple peers,or be reused as part of a larger RCF function. With the RCF functions utilizing standard programming expressions (AND, OR, NOT etc.) and familiar software constructs, a simpler and repeatable programmable approach to routing policies can now be defined at scale.

Scalable traffic policies

Routing deployments at the network edge for Internet peering and VPN services, require packet classification policies for network security, Policy Based Routing (PBR) and Quality of Service. On traditional routing platforms the implementation method of defining multiple discrete policies is resource intensive and limited by finite TCAM space, which therefore limits the scale and complexity of the policies that can be deployed. Arista's EOS software and the R-Series platforms implement a unique algorithmic model for defining packet classification policies, allowing overlapping fields and packet attributes to be shared across policies, enabling complex multifield classifiers to be expressed in fewer TCAM entries, thereby providing a simpler model for defining policy while greatly improving scale by optimizing TCAM resources.

FlexRoute™

The innovative FlexRoute Engine of Arista's EOS software provides internet scale routing with the R-Series platforms, allowing the platforms to support 2.5M+ FIB routes and 30M+* RIB routes. The FlexRoute Engine is a patented algorithmic approach to constructing layer 3 forwarding tables on the R-Series platforms and is a key enabler to building scalable internet peering and VPN services. When compared to legacy LPM approaches, FlexRoute uses less active silicon (lower activity factor), combined with a more efficient use of chip resources to hold the LPM forwarding tables.

FlexServices

As an MPLS service edge router the R-Series platform provides support for traditional E-LINE (Pseudowire) , E-LAN (VPLS) and L3 VPNs (RFC 4364) services, along with next generation EVPN-MPLS and EVPN-VXLAN services. Support for VPWS, E-Tree, L2VPN and L3VPNs with seamless VPN stitching eases migration when evolving between the different E-LAN and L3VPN service models. EOS's innovative software approach to service definition, also means the MPLS transport is transparent to the VPN service type or model, thereby providing a fully flexible and user configurable preference based tunnel (SR-TE, RSVP-TE, LDP, SR, BGP-LU etc) route resolution.

FlexEncap

The innovative FlexEncap framework of Arista's EOS software on the R-Series platform is a new approach to addressing the complex packet manipulation requirements when delivering Layer 2 and L3 VPN services. FlexEncap allows QoS policy, packet transparency (protocol tunnelling) and VLAN tag manipulation action (consume, retain, re-write.) can be defined on a per service basis, with the packet manipulation automatically applied bidirectionally to ensure symmetric forwarding behavior.

Traffic engineering and Fast Reroute

As an MPLS LER and LSR node in existing core networks, the R-Series platform, in combination with Arista' EOS software provides support for RSVP-TE for traffic engineering and Fast ReRoute (FRR). For next-generation software-defined MPLS core solutions, the platforms offer a rich MPLS Segment-Routing (SR) stack, with SR-TE support for traffic engineering requirements and TI-LFA for Fast ReRoute. EOS's flexible approach to label allocation and preference based tunnel (SR-TE, RSVP-TE, LDP, SR, BGP-LU etc) route resolution, also means the different TE models are not mutually exclusive and can be run concurrently to provide a seamless migration path between the different models when required.

Encryption

Data integrity is at the forefront of building the next generation WAN, but implementing a cost effective solution when data encryption is now required at scale over 100G/400G links has become a major challenge for today's Service Providers. Arista's R-Series platforms address this challenge by providing hardware based line-rate 100G/400G MACSec and IPSec encryption solutions. IEEE802.1AE MACSec solution provides point-to-point AES-256-GCM encryption typically deployed over dark fibre connections between sites, alternatively, IPSec may be deployed to deliver point-to-multipoint encryption across a routed infrastructure.

Arista R-Series Router Feature support

The tables below provide an overview of the EOS software feature support (for both IPv4 & IPv6 where applicable) across the R-Series platform; some features have a hardware platform dependency. For the official listing of feature support on specific hardware platforms and the relevant software release, please reference the Feature Support Matrix - https://www.arista.com/en/support/product-documentation/supported-features

	Features	
Layer 3	• L3 Sub-interfaces • Routing Protocols: Static Routes, OSPF, OSPFv3, BGP, MP-BGP, IS-IS, and RIPv2 • ISIS Multi-topology (MT) • 512-way Equal Cost Multipath Routing (ECMP) • UnEqual Cost Multipath Routing with BGP communities • VRF and VRF route-leaking agent • Bi-Directional Forwarding Detection (BFD) • Micro BFD (RFC 7130) • Unicast Reverse Path Forwarding (uRPF) • Prefix Independent Convergence (PIC)	• VRRP • Virtual ARP (VARP) • Policy Based Routing (PBR) • Route-Target Pruning • Route-Target constraint • RPKI • Route Maps • Route control functions (RCF) • Route Reflector • Optimal Route Reflector (ORR) • BGP Monitoring Protocol (BMP)
Multicast	• IGMP v2/v3 and MLDv2 • Protocol Independent Multicast (PIM-SM / PIM-SSM) • PIM-BiDir • Anycast RP (RFC 4610) • Multicast Source Discovery Protocol (MSDP)	• Multicast Only Fast ReRoute (MoFRR) • Multicast VPN (NG-MVPN) mLDP and Default MDT • Multicast VPN (NG-MVPN) mLDP and Data MDT • Multicast VPN (NG-MVPN) with RSVP-TE*

	Features	
Layer 2	• 802.1w Rapid Spanning Tree • 802.1s Multiple Spanning Tree Protocol • Rapid Per VLAN Spanning Tree (RPVST+) • 4096 VLANs • Private VLANs • Q-in-Q (IEEE 802.1ad) • Layer 2 sub-interfaces • Flexible VLAN encapsulation • 802.3ad Link Aggregation/LACP » 256 Ports / Channel » 2048 groups per system (subject to system density)	• MLAG (Multi-Chassis Link Aggregation) » Uses IEEE 802.3ad LACP » 512 ports per MLAG • 802.1Q VLANs/Trunking • 802.1AB Link Layer Discovery Protocol • 802.3x Flow Control • Jumbo Frames (9216 Bytes) • IGMP v1/v2/v3 snooping • Storm Control • Loop protection
EVPN	• EVPN with VXLAN forwarding plane • EVPN with MPLS forwarding plane • EVPN VLAN based services • EVPN VLAN-aware services • EVPN Multihoming All-Active/Active-Standby • EVPN L2 VPNs (Type-2 MAC routes) • EVPN VPWS with MPLS encap • EVPN VPWS-FXC with MPLS encap • EVPN E-tree RFC 8317 and RFC8317bis (VXAN encap) EVPN Integrated Routing and Bridging (IRB)	• EVPN IRB with Anycast GW • EVPN L3 VPNs (Type-5 IP prefixes) • EVPN L2 Multicast (VXLAN nd MPLS* encap) • EVPN OISM (VXLAN nd MPLS* encap) • EVPN OISM with PIM Edge Gateway (PEG) • EVPN-VXLAN/MPLS to IP-VPN GW • EVPN-VXLAN to EVPN-VXLAN L2/L3 GW • EVPN-VXLAN to EVPN-MPLS L2/L3 GW • EVPN-VXLAN to EVPN-VXLAN OISM GW*
L2 and L3 VPN	• IP-VPN (RFC 4364) • 6PE and 6vPE • Inter-AS Option A, B and C • iBGP for PE to CE interconnect (RFC 6368) • EVPN-VXLAN/MPLS to IP-VPN GW	• LDP signaled Pseudowires • Type-4 and 5 Pseudowires • LDP signaled VPLS • LDP signaled VPLS with BGP AD • BGP signaled VPLS with BGP AD*
Quality of Service (QoS)	• Up to 8 queues per port/sub-interface • WRR and strict priority queueing • 802.1p based classification • DSCP based classification and remarking • Egress shaping / Weighted round robin (WRR) • Policing per sub-interface	• Hierarchical shaping per sub-interface • Explicit Congestion Notification (ECN) marking • 802.1Qbb Per-Priority Flow Control (PFC) • 802.1Qaz Enhanced Transmission Selection (ETS) • Data Center Bridging Extensions (DCBX)
OAM	• Continuity Fault Management (MIP Support) • Continuity Fault Management (UP MEP & DOWN MEP support) • CFM Loss Measurement tests (LM) • CFM Synthetic Loss Measurement (SLM) • CFM Delay Measurement (DM)	• Link fault Signaling • RFC 2544 (initiator & reflector) • TWAMP (extension) • EOS Connectivity Monitor • MPLS Ping and Traceroute • VCCV Support

	Features	
Security	• Control Plane Protection (CoPP) • Ingress / Egress ACLs using L2, L3, L4 fields • Ingress / Egress ACL Logging and Counters • Scalable traffic policies • MAC ACLs • ACL Deny Logging • ACL Counters • Atomic ACL Hitless restart • DHCP Relay / Snooping	• MAC Security • TACACS+ • RADIUS • ARP trapping and rate limiting • MACsec • IPSec • BGP FlowSpec • VXLANSec
Label Distribution	• LDP and MLDP (LER/LSR) • ISIS-SR • OSPF-SR • Node-SID, Proxy-SIDI, Anycast-SID, Prefix-SID, Binding SID • Adjacency SID (local, Global & static) • BGP-LU • BGP-SR	• RSVP-TE • BGP-LU for EPE • TI-LFA • Configurable AF tunnel-rib resolution • Flow-Aware Transport (FAT) Label, RFC 6391 • Entropy label RFC 6790 • MPLS over GRE*
Traffic Engineering (RSVP-TE)	• RSVP-TE (LER/LSR) with GR • Explicit and Dynamic (strict/loose) paths • Extended Administration groups • Secondary LSP, Cold and Hot standby • Bandwidth Reservation • Auto-bandwidth reservation • Adaptive split tunneling • IGP shortcuts • Path prioritization	• Path re-optimisation • Tunnel profiles • IPv6/IPv4, 6PE and VPN (IP-VPN, VPWS, L2 EVPN, L3 EVPN) steering into RSVP-TE policies • Class Based Forwarding (CBF) into RSVP-TE Tunnel • Fast ReRoute (FRR) per tunnel • Colored tunnel-rib resolution • LDP over RSVP-TE
Traffic Engineering (SR-TE)	• SR-TE Endpoint colored based policies • On-box and off-box API policy provisioning • Dynamic SR-TE policies • Prioritization of candidate paths within policy • Load-balancing across candidate paths based on weight • Seamless BFD with Round Trip Time (sBFD with RTT) • Colored Route-map for steering into SR--TE policy	• IPv6/IPv4, 6PE and VPN (IP-VPN, VPWS, L2 EVPN, L3 EVPN) steering into SR--TE policies • Class Based Forwarding (CBF) into SR--TE Tunnel • BGP-LS • ISIS over GRE • TWAMP-lite for delay metric calculation • ISIS FlexAlgo (TE metrics, dynamic delay*) • ISIS FlexAlgo with TI-LFA • Configurable AF tunnel-rib resolution

	Features	
Network Management	• CloudVision • Configuration rollback and commit • 100/1000 Management Port • RS-232 Serial Console Port • USB Port • SNMP v1, v2, v3 • Management over IPv6 • Telnet and SSHv2 • Syslog • AAA	• Industry Standard CLI • Beacon LED for system identification • System Logging • Environment monitoring • Arista eAPI (http & https) • NETCONF • RESTCONF • gNMI, gNOI, gNSI, gRIBI • OpenConfig data models • EOS SDK
Advanced Monitoring and Provisioning	• Latency Analyzer and Microburst Detection (LANZ) • Zero Touch Provisioning (ZTP) » Advanced Mirroring » Port Mirroring » Enhanced Remote Port Mirroring » GRE enabled SPAN » SPAN/TAP M:N Aggregation » L2/3/4 Filtering • Advanced Event Management suite (AEM) » CLI Scheduler » Event Manager » Event Monitor » Linux tools with shell access	• Integrated packet capture/analysis with TCPDump • Restore and Configure from USB • RFC 3176 sFlow • IPFIX • Optional SSD for logging and data capture • IEEE 1588 PTP • SyncE • Postcard telemetry

** Supported in a future EOS release*

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, 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