



AVAYA

The Power of We™

Avaya Ethernet Routing Switch 8800

Resilient, flexible, and scalable, the Avaya Ethernet Routing Switch 8800 delivers full-featured network virtualization, exceptional value and cost-effectiveness.

Virtualization can transform your IT infrastructure and your business by providing a clear path to advanced applications that create a distinct competitive advantage. Virtualization delivers flexibility and scalability, and enables faster activation of new services in data centers and across the campus. Leveraging high-availability and high-performance, virtualizing servers and consolidating services delivers benefits such as accelerated decision making, decreased recurring costs, increased productivity and simplified management.

The Avaya ERS 8800 offers exciting new possibilities with the Avaya VENA Fabric Connect capability, based on the Shortest Path Bridging (SPB) standard from the IEEE. Leveraging the power of this standardized technology, together with unique, value-added extensions, Avaya delivers a solution that is genuinely enterprise-ready and truly optimized. The Fabric Connect solution extends beyond the data center into the core of the campus: it offers simplified configuration and management, including efficient service activation free of error and delay, optimized traffic separation that enables multi-tenant partitioning and regulatory compliance, and delivers the industry's only optimized end-to-end Cloud architecture. Avaya VENA Fabric Connect creates a new approach to data networking, one that helps enterprises fully reap the benefits of virtualization, in a more simplified and cost-effective manner. This technology utilizes a new, end-to-end enterprise-wide architecture to help CIOs and IT departments meet the surging demand for new content and business collaboration applications.

Highlights of the Ethernet Routing Switch 8800

- Innovates with Avaya VENA Fabric Connect; delivering game-changing Layer 2 and Layer 3 virtualization capabilities for the data center and beyond

- Simplified configuration & management including efficient service activation - free of error and delay
- Optimized traffic separation - helping to ensure multi-tenant partitioning & regulatory compliance
- Delivers the only optimized end-to-end cloud architecture
- Offers the Industry's leading resiliency model - Avaya Switch Clustering - empowering the most demanding applications and boosting performance by forwarding Layer 2 & 3 traffic across all available links
- Features unique field-reprogrammable NPU-based Interface Modules that, unlike conventional ASIC-based hardware, maintain full hardware-based performance and optimization as functionality and services evolve
- Enables flexible virtualized Layer 3 deployment scenarios with device and network options: VRF-Lite, the innovative Avaya IP VPN-Lite, MPLS, & IETF IP VPN
- Enables consistent IP VPN services delivered across the campus and metro; leveraging the same infrastructure to seamlessly extend service provider MPLS-networks into the LAN
- Supports high-performance IPv6 networking - a key scalability tool for demanding and expanding networks
- Offers high-density 10G, very high-density Gigabit and 10/100/1000 Ethernet for enterprise core and aggregation applications, delivering competitively-high value, flexibility, and enhanced slot conservation with a Combo option
- Best-in-class Switch Clustering resiliency model is extended to VMware Server virtualization in an iSCSI storage area network environment

Companies turn to technology to help boost the bottom line and to increase productivity, advances in one area of technology can often lead to real challenges in others. Virtualization is a case in point, particularly when it comes to efficiently connecting a myriad of disparate applications and systems - many now virtualized - across multiple locations.

- Supported by a unified management framework featuring consistent AJAX-compliant web-based common services, authentication and audit logging, also benchmarks network traffic and identifies anomalous behavior using Standards-based IP Flow Information Export (IPFIX)

The Ethernet Routing Switch 8800 meets demanding enterprise-class requirements for scalability, simplification, maximized application uptime, value, and security. It reduces network design complexity by simplifying the network architecture and increasing per port value with advanced features on high-density modules.

In addition to the capabilities offered by Avaya Fabric Connect, multiple options are available with the ERS 8800 for creating IP Virtual Private Networking solutions across the entire enterprise. Avaya Layer 3 virtualization is simple, flexible, and easy to deploy – and doesn't require adjustments to your existing infrastructure, requiring no additional capital equipment expenditures. Because solutions are standards-based and use well-understood IP techniques, less training time is required. This can reduce operational costs when compared to the knowledge transfer requirement of rival offerings.

Providing leading Ethernet port densities, the Avaya ERS 8800 delivers the infrastructure for a highly reliable network that empowers unified communications and other business-critical applications.

Key enabling technologies for improved business collaborative include highly-effective unified communications; crucial to a successful deployment is the selection and implementation of a reliable and versatile infrastructure that will help ensure constant availability.

New for the 7.2 release:

- Introducing support for the Avaya VENA Unified Access solution; the ERS 8800 delivers the data path optimization component

- Enhancing the Avaya VENA Fabric Connect capability with the introduction of the Multicast-over-SPB feature. This feature concurrently supports three deployment scenarios:

- L2 Virtual Service Networks with Multicast
- L3 Virtual Service Networks with Multicast
- IP Shortcuts with Multicast

Avaya VENA Unified Access

Balancing cost, performance, and accessibility is a case of perpetual tradeoff in today's WLANs – none of the traditional architectural approaches provides a satisfactory resolution. A new solution is needed, one that leverages the power of the centralized WLAN Controller without sacrificing performance, but also allowing for cost-effective distribution of WLAN Access Points. The answer is Avaya VENA Unified Access.

Instead of focusing solely on whether the data plane is best handled in the WLAN Access Point or the WLAN Controller, the Unified Access approach integrates the wireless data plane directly into the wired networking infrastructure – completely decoupled from the classic WLAN components. In effect, this makes the hardware burden for switching WLAN traffic simply disappear.

Here's how: traditional Ethernet Switches already have most of the capabilities to handle WLAN traffic natively – the issue is not processing power, but a lack of awareness of the mobility context for roaming sessions; that is, when a mobile device that has moved around the campus and therefore has an IP Address from a new subnet. To facilitate efficient coordination with the WLAN control plane, and a lightweight mobility agent added to the switch (this will program switching tables in real-time as sessions move), this awareness is easily incorporated. It makes much more sense to leverage the power of the wired network, rather than hamstringing the WLAN by inserting software switching points through out, or backhauling it all to another location just to make switching decisions.

Delivering ubiquitous access to enterprise applications and unified communications tools can have dramatic impact on enterprise productivity and performance. Avaya is delivering this with its Unified Access solution. By integrating the wireless data plane with the wired data network – specifically the ERS 8800 – the solution overcomes the disadvantages of both the centralized and distributed approaches and creates a scalable, reliable, high-performance architecture; perfect for 'Era of Mobility'.

The Avaya VENA Unified Access solution delivers:

- Genuine, next-generation scalability and performance
- Real high-availability and robustness
- A consistent access control solution
- Lower equipment and operational costs
- Enables virtualization of the control plane
- Unifies wired and wireless management

Unified Access delivers the intelligent edge, one that helps reduce the enterprise's total costs through the efficiencies of a network that is smart enough to authenticate, auto-provision, and provide security, all without constant intervention by the network administrator. The future is now, Avaya is delivering this today with its Unified Access and Identity Engines authenticated network access solutions.

Multicast enhancements to Avaya VENA Fabric Connect

Many of today's networks are constrained when attempting to make use of applications that require Multicast. The problems include inefficient bridged IP Multicast configurations, or routed IP Multicast that requires complex protocols; making networks burdensome to configure and operate. With the 7.2 software release, the ERS 8800 builds on the simplicity and agility introduced with the Avaya VENA Fabric Connect capability, by leveraging extensions to the Shortest Path Bridging (SPB) technology, that support a truly optimized Multicast solution.

Both Bridged and Routed IP Multicast traffic are supported, avoiding the inefficiencies or complexities that exist today, without any additional technology or protocols. This optimized functionality is achieved by the straightforward expediency of extending SPB's control plane to exchange advertisement and membership information for IP Multicast streams. This innovative means that Fabric Connect now provides a complete solution for Layer 2 Unicast, Broadcast, and Multicast virtualization, as well as Layer 3 Unicast and Multicast routing and forwarding virtualization.

The Multicast over SPB feature supports three operational modes:

- L2 VSN with Multicast – Virtual Services Networks with IGMP configured on access networks delivering optimized forwarding of Multicast traffic in a bridged environment
- L3 VSN with Multicast – Virtual Services Network with VRF-based Routed IP Multicast configured within the Fabric and IGMP configured on access networks
- IP Shortcuts with Multicast – Routed IP Multicast leveraging the Global Routing Table for direct, Layer 3 reachability within the Fabric and IGMP configured on access networks

The benefits of this innovation center around the simplicity of configuration and the optimization of how Multicast-based applications operate. The absence of a Root Bridge requirement – with no election process and no blocked links – is itself a major step forward in terms of efficiency and availability. Added to this, leveraging further extensions to the SPB model enable Avaya to deliver a solution that is agile, flexible, simple, stable, highly scalable, and notable for the ability to be quickly

deployed; time-to-service being the catchword for modern businesses.

What is Avaya VENA?

- Avaya Virtual Enterprise Network Architecture (Avaya VENA) allows organizations to more easily optimize business applications and service deployments in and between data centers and campuses
- Avaya VENA helps reduce costs and can significantly improve time to service by more efficiently optimizing the network connections between application servers and end users – essentially providing an end-to-end connection from the desktop all the way to the data center
- It can also reduce risks for CIOs by simplifying provisioning and policy configurations, reducing network re-designs, and providing new, streamlined tools for network management
- It can simplify device and network configuration; replaces multiple existing technologies with a single unified technology
- It can optimize connectivity and resource availability/utilization; is not constrained by physical topology and works equally well with mesh or ring installations, small or large
- It can increase stability and scalability; decoupling the network core from user/application transport providing a completely robust solution
- It pushes configuration out to the edge of the cloud, automates connectivity, and can dramatically improve time-to-service; can remove any need to re-configure the core, removing the risk of human error, and it also empowers provisioning on the fly

- It can seamlessly expand from the data center – where most demand exists today – out to the greater campus and beyond to the broader enterprise; delivering streamlined end-to-end connectivity from users to their content, all using one consistent technology and design model

To find out more about Avaya VENA please visit us at:

www.avaya.com/vena

Resiliency, intelligence and scalability without design complexity

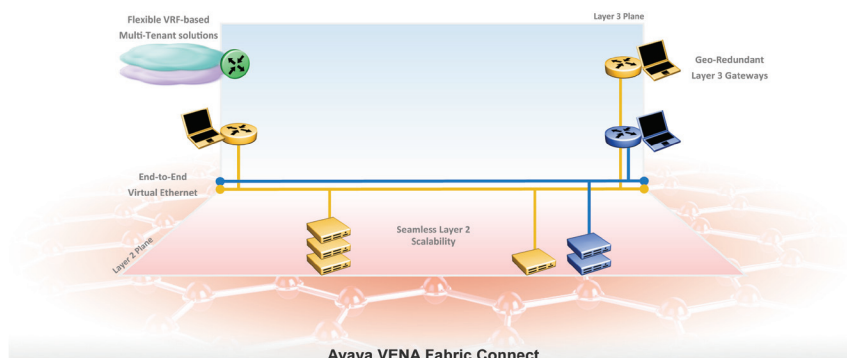
The Avaya ERS 8800 is a proven, tested, resilient, and intelligent network solution that scales, delivering to the core of the network hundreds of Gigabits per second (Gbps) and hundreds of millions of packets per second (Mpps) of real-world performance. Its flexible architecture reduces the complexity of network design, making it ideal for large-scale enterprise campuses.

The ERS 8800 is a balanced solution, unconstrained by bottlenecks imposed by inferior designs. In addition to establishing a solid foundation for traditional unified communications, the ERS 8800 delivers a flexible networking infrastructure that fosters growth by enabling businesses to leverage new, emerging applications and technologies with a unique architecture which provides optimum performance.

Always-on availability

Network resiliency is a basic mandatory requirement when implementing a modern unified network. The ERS 8800 supports redundant connectivity for virtualized solutions with technologies such as Avaya VENA Fabric Connect, VRF-Lite, VPN-Lite, and MPLS LER IP-VPN for Edge networks.

Avaya Fabric Connect is built on enhanced IEEE Shortest Path Bridging to provide resiliency, simplicity and a consistent interconnect that transparently supports co-existing services. Fabric Connect spans the network and enables one-touch provisioning for a full range of capabilities provided by



Real-world scenario: the NPU advantage

A well-known rival product – one of the many that utilize traditional ASIC technology – is only able to deliver IPv6 forwarding performance at 50% of the levels claimed for IPv4. This is a prime example of how newer features cannot be guaranteed at the same hardware-based performance levels unless there is an ability to upgrade the architecture; this is the advantage that the Avaya unique NPU design provides, and we continue to leverage this for evolving and emerging functionality such as network virtualization and Shortest Path Bridging-based Fabric Connect.

multiple Virtual Services Networks (VSNs). This architecture can increase flexibility and scalability by delivering an infrastructure that creates a 'private cloud', to deliver always-on content and access to applications in a dramatically simplified model. This approach also can protect enterprises' core networks from the costly failures and human-error issues that are often experienced by the traditional, complicated process of provisioning applications in a virtualized environment.

To provide maximum protection, the ERS 8800 addresses resiliency at multiple levels. At the hardware level, the switch provides hot-swappable modules and fan trays along with N+1 and dual input power supplies. Its software delivers resiliency for the core with Industry-leading features that include Virtual Link Aggregation Control Protocol (VLACP) for Layer 1-2 link failure detection, Bi-Directionally Forwarding Detection (BFD) for Layer 3 link failure detection, and Switching Clustering that leverages our pioneering Split Multi-Link Trunking (SMLT), Routed Split Multi-Link Trunking (R-SMLT), and VRRP Active/Active technologies.

Additionally, organizations are encouraged to dual-connect servers and, with minimum additional investment, the sub-second failover advantage is automatically extended beyond the boundary of the networking equipment, all the way to the application host. Competitive solutions, basing their failure recovery model on variations of the Spanning Tree Protocol, cannot provide a comparable level of resiliency and simplicity.

Future-proofing the network


Network devices must be able to distinguish different traffic types and to handle different traffic requirements. A sense of traffic class awareness combined with the ability to process each type uniquely sets the intelligent network apart from common-place offerings. The ERS 8800 combines intelligence and performance to create a next-generation intelligent network solution.

The networking industry is a perpetual work-in-progress, an 'unfinished masterpiece', and the number of standards now runs into the thousands. Equipment that is based on a ASIC traditional architecture is limited in that the functionality of these chipsets are set at a certain point in time and cannot easily adapt to future change. Typically this means that newer features and functionality are no longer supported only in hardware and require additional software processing. The switching architecture of the ERS 8800 is uniquely based on Network Processing Units (NPU) rather than the ASIC technology typically found in rival products. NPUs are large-scale CPU arrays specifically designed for network-related functions such as efficient examination and manipulation of packet headers. The specialized high-performance NPU of Avaya is known as the Route Switch Processor (RSP) and is an in-house development. It delivers fast-path, protected for the lifecycle of the product, through its ability to support in-life firmware upgrades and provides 10Gbps line rate switching and routing capabilities regardless of standards evolution. Avaya has been able to leverage this re-programmable capability to deliver new innovative features, such as IPv6, our flexible IP VPN suite, and the new Fabric Connect capabilities, and is unique in being able to help ensure emerging functionality is continuously delivered with hardware-based performance levels.

The ERS 8800 reduces complexity and risk in network design by simplifying the network architecture and increasing value with advanced features on high-density modules. High port density, combined with rich capabilities and leading reliability




ERS 8800 3-Slot Switch



ERS 8800 6-Slot Switch



RSP-based Interface Modules



ERS 8800 - 10-Slot Switch

technologies, deliver exceptionally high value to the enterprise. Avaya RSP technology, based on this flexible NPU architecture, offers investment protection with an in-field firmware upgrade capability, and helps ensure that the ERS 8800 remains ever-green, always delivering hardware-based performance.

Innovative options

Flexible and Scalable

The ERS 8800 is available in multiple Chassis options: a 10-Slot featuring eight Slots for Interface Modules; a 6-Slot featuring four Slots for Interface Modules; and a 3-Slot featuring two Slots for Interface Modules. This enables the ERS 8800 to be deployed in a variety of deployment scenarios, always helping to ensure the optimum mix of flexibility, capacity, and cost-effectiveness. There is also the 8010co Central Office – 10-Slot Chassis which is NEBS-compliant for deployments scenarios that require a carrier-class platform.

The 8895SF Switch Fabric/CPU Module is the latest version and offers significant enhancements in terms of CPU performance and Memory capacity; it is also 50% more energy-efficient. These advances allow the 8895SF to natively support both existing and emerging services and applications that place intense demands on the networking infrastructure. The 8895SF is functionally equivalent to the previous 8692SF Switch Fabric/CPU Module when upgraded with the Enterprise Enhanced CPU Daughterboard (aka SuperMezz).

Versatile Interface Options

The ERS 8800 supports an innovative hybrid combination module that concurrently supports 10G Ethernet (x2), 100BASE-X (x24), and 10/100/1000 (x8) ports; economical, flexible and a class-first, this ‘combo’ module meets the demands of smaller aggregation sites. It is an affordable solution providing all the functionality many enterprises need in one convenient and cost-effective module.

The range of high-performance I/O module options offer a number of practical options, and these include a high-density 10G Ethernet (12 ports per module and up to 96 ports per Chassis) and high-density 48-port modules support 10/100/1000 Copper or

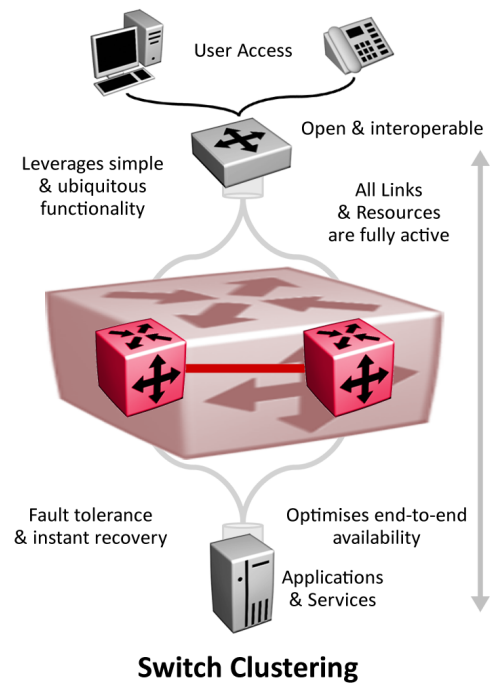
1000BASE-X SFP connectivity. These latest generation 8800-Series Interface Modules featuring the RSP 2.7 network processing units, enable the ERS 8800 to deliver enhanced mirroring capabilities, including one-to-many, many-to-one and many-to-many mirroring for sophisticated traffic analysis and IDS/TPS clustering.

Applications: new capabilities & choices optimize communications

Through advanced capabilities including Avaya Fabric Connect, VRF-Lite, IP VPN-Lite, and MPLS-based IP VPN, the ERS 8800 enables virtual services across campus and metro environments to address new business requirements and applications. For example, a university seeking to connect several campuses can choose traditional MPLS technology or can leverage the innovative Avaya IP VPN-Lite solution, and where an Ethernet-based MAN is available, transition directly to enhanced and extended SPB.

SPB is based on an open standard, 802.1aq developed by the IEEE (Institute of Electrical and Electronic Engineers). Being both standardized and open this means that as other vendors also introduce support for SPB there is a genuine opportunity for interworking between different devices from different suppliers; empowering customers with more freedom. SPB has been described as “Layer 2 Routing”, and there is a developing requirement to operate much of the data center network at Layer 2 (as opposed to Layer 3).

As virtualization becomes more widely deployed, networks will inherently become much more dynamic than they are today, and having an end-to-end solution that simplifies the perpetual adds, moves and changes will become crucial – especially given the fact that human error in provisioning new services accounts for almost 40% of all network downtime. This is just as important in the campus network as it is in the data center; Avaya, like our customers, understands that availability is an end-to-end equation.



With the Avaya VENA Fabric Connect capability, virtual networks can be created in seconds with single touch provisioning. The core is built out and provisioned one time and when new virtual networks are required - either to support new communities of interest or new applications - they are enabled solely at the edge of the network without requiring any manual provisioning of the core. This reduces service activation from the days and months, necessitated by today's legacy networks, to minutes and seconds. It essentially eliminates the human error factor and thereby dramatically increases network availability.

In many organizations, the network requires a total separation of traffic. This could be for regulatory or security reasons or it could be because a central IT organization has to support departments or subsidiary companies, each with their own unique needs, however at the same time all requiring access to a common set of business applications or services. With the extension of the Avaya VENA Fabric Connect capability into the campus network, setting up each of these distinct communities of interest is both simplified and optimized. Now, traffic separation can be truly end-to-end, without traffic being re-combined into shared networks

within the core; a very common limitation of legacy networks.

Avaya is the vendor creating an enterprise-wide model for delivering highly resilient access to applications and services with light-touch provisioning. This ultimately reduces operational expenses and simplifies the overall management of the infrastructure. Avaya has defined a single consistent model that spans from user to content, one that retains the most efficient elements of our proven always-on heritage together with the advantages of a sophisticated, standards-based and carrier-proven capability. Crucially, Avaya Fabric Connect does not suffer from the typically limitations of single points-of-failure and the need to interconnect differing forms of technology, one for the data center and others for the campus and MAN/WAN.

Virtual Routing and Forwarding (VRF-Lite)

Through VRF-Lite on ERS 8800, enterprises can use the same hardware platform to create multiple Layer 3 routing domains to support multiple customers and to maintain separation for both Unicast and Multicast traffic.

The VRF-Lite capability virtualizes routing within the switch, addressing business and networking challenges driven by activities such as mergers and acquisitions, data center consolidation, departmental or business unit segmentation, and evolving audit and compliance requirements. By enabling the switch to have multiple routing instances (up to 255), more sophisticated connections can be enabled in addition to support for overlapping IP Addresses. Complete and total traffic separation at Layers 2 and 3 is the usual practice, however the system can be configured to provide inter-VRF forwarding capabilities, allowing shared access to common resources.

Virtual Private Networking through IP VPN-Lite

The Avaya IP VPN-Lite capability is an innovative IP-in-IP technology that delivers VPN services that are easier to implement, deploy, and manage. With IP VPN-Lite, enterprises can

Real-world scenario: Avaya Fabric Connect

“Metro City” has a large number of departments – each with their own systems and applications – but each requiring access to some shared resources. Due to the sensitive nature of some of the traffic within these departments, providing complete isolation of traffic is of paramount importance. Because of budget constraints, “Metro City” was looking for a cost-effective solution that would be easy to manage and maintain in order to keep on-going operational costs low. As one of their primary cost reduction initiatives they are looking at server virtualization in their data center, and at potentially mobilizing virtual machines in the future.

“Metro City” chose to evolve their network with the Avaya ERS 8800, with which they can very easily establish distinct virtual networks to support each of their different departments. Since some of these departments carry sensitive traffic they liked that Avaya VENA’s architecture provides end-to-end traffic isolation – and unlike today’s legacy implementations – traffic is not recombined into shared networks in the core. They also liked the flexibility in setting up these different virtual multi-tenant networks since they can be enabled either at Layer 2 or at Layer 3 (by mapping, respectively, VLANs or VRFs to service Instances for transport over the end-to-end Fabric Connect solution).

Since their migration to data center virtualization was to occur progressively, “Metro City” appreciated the fact that Avaya’s architecture could be enabled in parallel with all other protocols presently in use on the network and there would be no need to change physical connections or existing configurations; they could migrate at their own pace, in a non-disruptive manner.

build any-to-any private connections between local or geographically dispersed sites using any IP infrastructure (private networks or via a public IP Service Provider).

IP VPNs are typically used for cross-location connectivity and to create trusted connections to external partner organizations, leveraging IP as the common carriage and removing dependency on specific wide area technologies (such as

Frame Relay or ATM) or exclusivity to a single service provider. Avaya IP VPN-Lite solution is inherently less complex and therefore much more cost-effective than using a full MPLS alternative. Managing IP VPN-Lite versus MPLS is simpler and does not require specialized carrier-class IT skills or resources. The foundation of IP VPN-Lite is simply an IP network, using the flexible RFC 2547/4364 connectivity model and it does not require an MPLS-enabled core

infrastructure. This simplified solution can scale per carrier-class MPLS with the cost-effective simplicity of a solution designed specifically for the enterprise. Delivering total flexibility, the ERS 8800 supports classic MPLS in addition to IP VPN-Lite and VRF-Lite and all VPN technologies can be concurrently leveraged to deliver individually tailored solutions.

Multi-Protocol Label Switching (MPLS)

MPLS forms the basis for most service provider IP VPNs and is used in most WAN solutions because it delivers sophisticated connectivity and traffic engineering techniques. By implementing this same functionality the ERS 8800 can interoperate directly with third party MPLS networks and participate in their IP VPNs, extending them into the enterprise campus network as required. Enterprise architects can leverage this interoperability to create ERS 8800-based MPLS environments that increase the overall level of transparency.

Enterprise-class and enterprise-ready: fit-for-purpose

Multicast Virtualization

Multicast virtualization of IGMP and PIM-SM/SSM, along with Unicast virtualization, is supported on the same system by leveraging VRF-Lite. This extends the benefit of simplified network design, less investment in hardware, lower operational expense; all of which is supported by the best-in-class Switch Clustering resiliency.

Avaya Automatic QoS

With Avaya Automatic QoS, an ERS 8800 supporting an Avaya Unified Communications solution automatically recognizes the special, private Differentiated Service Code Point (DSCP) values used by specific Avaya VoIP applications and optimizes the management of egress queues. Without this functionality, operators would require detailed knowledge of how QoS works and the private DSCP values in order to manually configure optimized queue usage. With this feature, the process is automated, optimized, and protects against errors in configuration.

Real-world scenario: multi-tenant operations

An airport handling a large amount of traffic from multiple national and international airlines and local on-site businesses, seeks a networking solution that maximizes application uptime, safeguards information, and delivers excellent business value. By choosing the Ethernet Routing Switch 8800 with its versatile Fabric Connect networking capabilities, the airport supports all of these communities and keeps traffic segregated using a single, cost-effective hardware platform that is easy to implement and manage.

Management and Visibility

The ERS 8800 can be managed by a variety of management tools, creating a flexible operational environment based on business requirements. These include: standardized Command Line Interface (CLI), Web-based Enterprise Device Manager (EDM), SNMP-based management (SNMPv1, v2 & v3), and the Avaya Data Solutions Management framework provides for comprehensive, centralized, and multi-faceted network management. This framework is based on Common Services – authentication and access control, audit, etc – and then a number of integrated AJAX-based plug-in applets that deliver seamless task-specific capabilities all with a consistent look and feel: Configuration and Orchestration Management; Visualization, Performance, and Fault Management; IP Flow Manager; and the Virtualization Provisioning Service.

Provision Wizards and other labor-saving tools provide faster service activation and more consistent approach to configuration; this has the added benefit of reducing human-error, as templates are pre-populated with best-practice recommendations or required values. The entire framework is context-based which enables a faster, more accurate and highly-intelligent approach to delivering both device-centric and network-wide management services.

The ERS 8800 also supports enhanced system-wide troubleshooting information providing comprehensive information should the CPU need to recover automatically. The Key Health Indicators (KHI) capability allows for the collection of statistics and information about system health for troubleshooting purposes, and identifies a small number of key health indicators that allow quick assessment of the overall operational state. Additional sophisticated system analysis tools such as RSP Packet Tracing and ERCD Records Dump enhance the serviceability and operational posture of the ERS 8800 platform.

Why choose Avaya?

Avaya offers multiple, flexible options that enable virtualized solutions across the campus including the standards-based Avaya Fabric Connect capability and the innovative IP VPN-Lite. The unique design architecture of the future-ready NPU-based Interface Modules sets Avaya apart; delivering optimal functionality and performance as new applications and services emerge, and offers high density and exceptional port value. The Combo Module – supporting Copper 10/100/1000, SFP, and XFP interfaces – cost-effectively meets the requirements of smaller sites. Avaya is the solution vendor to offer best-in-class resiliency for Unicast, Multicast, Virtualized, and IPv6 environments.

Summary

The Ethernet Routing Switch 8800 is a resilient, efficient, scalable solution that enables enterprises to build a truly unified communication-ready network infrastructure and provide reliable business continuity for critical applications; enterprises can scale converged and web applications network-wide with always-on resiliency. The ERS 8800 offers a high-performance architecture combining rich, advanced services for converged applications that enhance, protect and simplify network service and operations. Customers making strategic investments in their campus LAN infrastructure can rely on the ERS 8800 to create flexible solutions that match their business evolution. A provider of end-to-end solutions spanning voice, data, applications and network management, Avaya has the proven expertise to help businesses enhance revenue potential, streamline business operations, increase productivity and gain competitive advantage.

Learn More

To learn more about the Avaya Ethernet Routing Switch 8800, contact your Avaya Account Manager, Avaya Authorized Partner or visit us at: www.avaya.com.

Specifications

General and Performance

- Switch architecture: 720Gbps gross throughput
- Switch Fabric performance: up to 512Gbps in an Active/Active configuration
- Frame forwarding rate: up to 380Mpps
- Frame length: 64 to 1518 Bytes (802.1Q Untagged), 64 to 1522 bytes (802.1Q Tagged)
- Jumbo Frame support: up to 9,000 Bytes (802.1Q Tagged)
- Multi-Link Trunks: up to 128 Groups, with 8 Links per Group
- VLANs: up to 4,000 Port/Protocol/802.1Q-based
- Multiple Spanning Tree Groups: up to 32

- MAC Address: up to 64k
- IP Interfaces: 1,972
- Dynamic ARP Entries: up to 32k
- VRRP Interfaces: up to 255
- IP Forwarding Table: 250k
- ECMP Routes: up to 5k
- RIP Instances: up to 64
- RIP Interfaces: up to 200
- RIP Routes: up to 10k
- OSPF Instances: up to 64
- OSPF Adjacencies: up to 80
- OSPF Routes: up to 50k
- BGP Peers: up to 250
- BGP Routes: up to 250k
- SPB C-VLANs: up to 1,500
- SPB IS-IS Adjacencies: up to 50
- SPB IP Routes for L3 VSN: up to 25k
- SPB IP Routes for IP VPN-Lite/SPB: up to 250k
- VRF-Lite Instances: up to 255
- MPLS LDP LSPs: up to 16k
- MPLS Tunnels: up to 2,500
- PIM Active Interfaces: up to 200
- PIM Neighbors: 80/up to 200 for all VRFs
- IP Multicast Streams: up to 4k

System

- 8010 or 8010co 10-Slot, 8006 6-Slot, and 8003R 3-Slot Chassis options
- 8895SF Switch Fabric/CPU Module

Interface Modules

- 8812XL 12-port 10GbE SFP+ Ethernet Interface Module*
- 8834XG 34-port Combo 10/100/100 & 1GbE SFP & 10GbE XFP Ethernet Interface Module
- 8848GB 48-port 1GbE SFP Ethernet Interface Module
- 8848GT 48-port 10/100/1000 Ethernet Interface Module

Software

- ERS 8800 Base Software License

- ERS 8800 Advanced Software License, activating additional features: BGP4, IPv6 Routing, BFD, MSDP, & PCAP
- ERS 8800 Premier Software License, activating additional features: Advanced features plus Fabric Connect (including Multicast), IP VPN-Lite, MP-BGP, MPLS IP VPN, & VRF-Lite (including Multicast)

IEEE and IETF Standards

Compatibility

- 802.1D-1998 Spanning Tree Protocol
- 802.1p Priority Queues
- 802.1Q Virtual LANs
- 802.1s Multiple Spanning Trees
- 802.1w Rapid Reconfiguration of Spanning Tree
- 802.1v VLAN Classification by Protocol & Port
- 802.1X Port Based Network Access Control
- 802.1ag OA&M compliant for Loopback and Link-trace
- 802.1aq Shortest Path Bridging/D2.5
- 802.1AX/802.3ad Link Aggregation Control Protocol
- 802.3 CSMA/CD Ethernet (ISO/IEC 8802-3)
- 802.3 10BASE-T Ethernet
- 802.3i 10BASE-T - Auto-Negotiation
- 802.3u 100BASE-TX Fast Ethernet (ISO/IEC 8802-3, Clause 25)
- 802.3u 100BASE-FX
- 802.3u Auto-Negotiation on Twisted Pair (ISO/IEC 8802-3, Clause 28)
- 802.3x Flow Control on the Gigabit Uplink port
- 802.3z Gigabit Ethernet 1000BASE-SX & LX
- 802.3ab 1000BASE-T Ethernet
- 802.3ab 1000BASE-LX Ethernet
- 802.3ab 1000BASE-ZX Ethernet
- 802.3ab 1000BASE-CWDM Ethernet
- 802.3ab 1000BASE-SX Ethernet

*The 8812XL Module supports SFP+ Pluggable Transceivers that operate at 10Gbps only, not those that operate at 1Gbps. The 8812XL does support Direct Attach Cables for low-cost copper-based 10Gbps connectivity.

- 802.3ab 1000BASE-XD Ethernet
- 802.3ab 1000BASE-BX Ethernet
- 802.3ae 10GBASE-X XFP
- RFC 768 UDP Protocol
- RFC 783 TFTP Protocol
- RFC 791 IP Protocol
- RFC 792 ICMP Protocol
- RFC 793 TCP Protocol
- RFC 826 ARP Protocol
- RFC 854 Telnet Protocol
- RFC 894 A standard for the Transmission of IP Datagrams over Ethernet
- Networks
- RFC 896 Congestion control in IP/TCP internetworks
- RFC 903 Reverse ARP Protocol
- RFC 906 Bootstrap loading using TFTP
- RFC 950 Internet Standard Sub-Netting Procedure
- RFC 951 / RFC 2131 BootP / DHCP
- RFC 1027 Using ARP to implement transparent subnet gateways/ Nortel Subnet based VLAN
- RFC 1058 RIPv1 Protocol
- RFC 1112 IGMPv1
- RFC 1253 OSPF
- RFC 1256 ICMP Router Discovery
- RFC 1305 Network Time Protocol v3 Specification, Implementation and Analysis3
- RFC 1332 The PPP Internet Protocol Control Protocol
- RFC 1340 Assigned Numbers
- RFC 1541 Dynamic Host Configuration Protocol
- RFC 1542 Clarifications and Extensions for the Bootstrap Protocol
- RFC 1583 OSPFv2
- RFC 1587 The OSPF NSSA Option
- RFC 1591 DNS Client
- RFC 1695 Definitions of Managed Objects for ATM Management v8.0 using SMIv2
- RFC 1723 RIP v2 – Carrying Additional Information
- RFC 1745 BGP / OSPF Interaction
- RFC 1771 / RFC 1772 BGP-4
- RFC 1812 Router Requirements
- RFC 1866 HTMLv2 Protocol
- RFC 1965 BGP-4 Confederations
- RFC 1966 BGP-4 Route Reflectors
- RFC 1998 An Application of the BGP Community Attribute in Multi-home Routing
- RFC 1997 BGP-4 Community Attributes
- RFC 2068 Hypertext Transfer Protocol
- RFC 2131 Dynamic Host Control Protocol
- RFC 2138 RADIUS Authentication
- RFC 2139 RADIUS Accounting
- RFC 2178 OSPF MD5 cryptographic authentication/ OSPFv2
- RFC 2205 Resource Reservation Protocol – v1 Functional Specification
- RFC 2210 The Use of RSVP with IETF Integrated Services
- RFC 2211 Specification of the Controlled-Load Network Element Service
- RFC 2236 IGMPv2 for snooping
- RFC 2270 BGP-4 Dedicated AS for sites/single provide
- RFC 2283 Multiprotocol Extensions for BGP-4
- RFC 2328 OSPFv2
- RFC 2338 VRRP: Virtual Redundancy Router Protocol
- RFC 2362 PIM-SM
- RFC 2385 BGP-4 MD5 authentication
- RFC 2439 BGP-4 Route Flap Dampening
- RFC 2453 RIPv2 Protocol
- RFC 2475 An Architecture for Differentiated Service
- RFC 2547 BGP/MPLS VPNs
- RFC 2597 Assured Forwarding PHB Group
- RFC 2598 An Expedited Forwarding PHB
- RFC 2702 Requirements for Traffic Engineering Over MPLS
- RFC 2765 Stateless IP/ICMP Translation Algorithm
- RFC 2796 BGP Route Reflection – An Alternative to Full Mesh IBGP
- RFC 2819 Remote Monitoring
- RFC 2858 Multiprotocol Extensions for BGP-4
- RFC 2918 Route Refresh Capability for BGP-4
- RFC 2961 RSVP Refresh Overhead Reduction Extensions
- RFC 2992 Analysis of an Equal-Cost Multi-Path Algorithm
- RFC 3031 Multiprotocol Label Switching Architecture
- RFC 3032 MPLS Label Stack Encoding
- RFC 3036 LDP Specification
- RFC 3037 LDP Applicability
- RFC 3065 Autonomous System Confederations for BGP
- RFC 3210 Applicability Statement for Extensions to RSVP for
- RFC 3215 LDP State Machine
- RFC 3270 Multi-Protocol Label Switching Support of Differentiated Services
- RFC 3376 Internet Group Management Protocol, v3
- RFC 3392 Capabilities Advertisement with BGP-4 LSP-Tunnels
- RFC 3443 Time To Live Processing in Multi-Protocol Label Switching Networks
- RFC 3569 An overview of Source-Specific Multicast
- RFC 3917 Requirements for IP Flow Information Export
- RFC 4364 BGP/MPLS IP Virtual Private Networks
- RFC 4379 Detecting Multi-Protocol Label Switched Data Plane Failures
- draft-holbrook-idmr-igmpv3-ssm-02.txt IGMPv3 for SSM
- draft-ietf-bfd-v4v6-1hop-06 IETF draft Bi-Directional Forwarding Detection for IPv4 and IPv6 (Single Hop)
- RFC 1075 DVMRP Protocol

- RFC 1112 IGMP v1 for Routing/Snooping
- RFC 1519 Classless Inter-Domain Routing: an Address Assignment and Aggregation Strategy
- RFC 2236 IGMP v2 for routing / snooping
- RFC 2362 + some PIM-SM v2 extensions
- RFC 3446 Anycast Rendezvous Point mechanism using Protocol Independent Multicast and Multicast Source Discovery Protocol
- RFC 3618 Multicast Source Discovery Protocol
- RFC 3768 Virtual Router Redundancy Protocol
- RFC 1881 IPv6 Address Allocation Management
- RFC 1886 DNS Extensions to support IP version 6
- RFC 1887 An Architecture for IPv6 Unicast Address Allocation
- RFC 1981 Path MTU Discovery for IPv6
- RFC 2030 Simple Network Time Protocol v4 for IPv4, IPv6 & OSI
- RFC 2373 IPv6 Addressing Architecture
- RFC 2375 IPv6 Multicast Address Assignments
- RFC 2460 Internet Protocol, v6 Specification
- RFC 2461 Neighbor Discovery
- RFC 2462 IPv6 Stateless Address Auto-Configuration
- RFC 2463/4443 Internet Control Message Protocol for the Internet Protocol v6 Specification
- RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
- RFC 2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
- RFC 2526 Reserved IPv6 Subnet Anycast Addresses
- RFC 2710 Multicast Listener Discovery for IPv6
- RFC 2740 OSPF for IPv6
- RFC 2893 Configured Tunnels and Dual Stack Routing per port
- RFC 2893 Transition Mechanisms for IPv6 Hosts and Routers
- RFC 3056 Connection of IPv6 Domains via IPv4 Clouds
- RFC 3363 Representing Internet Protocol Version 6 Addresses in DNS3
- RFC 3484 Default Address Selection for IPv6
- RFC 3513 Internet Protocol Version 6 Addressing Architecture
- RFC 3587 IPv6 Global Unicast Address Format
- RFC 3596 DNS Extensions to Support IP v6
- RFC 3587 IPv6 Global Unicast Address Format
- RFC 3590 Source Address Selection for the Multicast Listener Discovery Protocol
- RFC 3596 DNS Extensions to support IP version 6
- RFC 3810 IPv6 Multicast capabilities SSH/SCP, Telnet, Ping, CLI, JDM support for IPv6
- RFC 1305 NTP Client / Unicast mode only
- RFC 1340 Assigned Numbers
- RFC 1350 The TFTP Protocol (Revision 2)
- RFC 2474 / RFC 2475 DiffServ Support
- RFC 2597 / RFC 2598 DiffServ per Hop Behavior
- RFC 1155 SMI
- RFC 1157 SNMP
- RFC 1215 Convention for defining traps for use with the SNMP
- RFC 1269 Definitions of Managed Objects for the Border Gateway Protocol v3
- RFC 1271 Remote Network Monitoring Management Information Base
- RFC 1304 Definitions of Managed Objects for the SIP Interface Type
- RFC 1354 IP Forwarding Table MIB
- RFC 1389 RIP v2 MIB Extensions
- RFC 1565 Network Services Monitoring MIB
- RFC 1757 / RFC 2819 RMON
- RFC 1907 SNMPv2
- RFC 1908 Coexistence between v1 & v2 of the Internet-standard Network Management Framework
- RFC 1930 Guidelines for creation, selection, and registration of an Autonomous System
- RFC 2571 An Architecture for Describing SNMP Management Frameworks
- RFC 2572 Message Processing and Dispatching for the Simple Network Management Protocol
- RFC 2573 SNMP Applications
- RFC 2574 user-based Security Model for v3 of the Simple Network Management Protocol
- RFC 2575 View-based Access Control Model for the Simple Network Management Protocol
- RFC 2576 Coexistence between v1, v2, & v3 of the Internet Standard Network Management Framework
- RFC 1212 Concise MIB definitions
- RFC 1213 TCP/IP Management Information Base
- RFC 1213 MIB II
- RFC 1354 IP Forwarding Table MIB
- RFC 1389 / RFC 1724 RIPv2 MIB extensions
- RFC 1398 Definitions of Managed Objects for the Ethernet-Like Interface Types
- RFC 1406 Definitions of Managed Objects for the DS1 and E1 Interface Types
- RFC 1414 Identification MIB
- RFC 1442 Structure of Management Information for version 2 of the Simple Network Management Protocol
- RFC 1447 Party MIB for v2 of the Simple Network Management Protocol bytes
- RFC 1450 Management Information Base for v2 of the Simple Network Management Protocol
- RFC 1472 The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol

- RFC 1483 Multiprotocol Encapsulation over ATM Adaptation Layer 5
- RFC 1493 Bridge MIB
- RFC 1525 Definitions of Managed Objects for Source Routing Bridges
- RFC 1565 Network Services Monitoring MIB
- RFC 1573 Interface MIB
- RFC 1643 Ethernet MIB
- RFC 1650 Definitions of Managed Objects for the Ethernet-like Interface Types using SMIv2
- RFC 1657 BGP-4 MIB using SMIv2
- RFC 1658 Definitions of Managed Objects for Character Stream Devices using SMIv2
- RFC 1695 Definitions of Managed Objects for ATM Management v8.0 using SMIv2
- RFC 1696 Modem Management Information Base using SMIv2
- RFC 1724 RIP v2 MIB Extension
- RFC 1850 OSPF MIB
- RFC 2021 RMON MIB using SMIv2
- RFC 2037 Entity MIB using SMIv2
- RFC 2096 IP Forwarding Table MIB
- RFC 2233 Interfaces Group MIB using SMIv2
- RFC 2452 IPv6 MIB: TCP MIB
- RFC 2454 IPv6 MIB: UDP MIB
- RFC 2465 IPv6 MIB: IPv6 General group and textual conventions
- RFC 2466 IPv6 MIB: ICMPv6 Group
- RFC 2578 Structure of Management Information v2
- RFC 2613 Remote Network Monitoring MIB Extensions for Switched Networks v1.0
- RFC 2665 Definitions of Managed Objects for the Ethernet-like Interface Types
- RFC 2668 Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units
- RFC 2674 Bridges with Traffic MIB
- RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
- RFC 2863 Interface Group MIB
- RFC 2925 Remote Ping, Traceroute & Lookup Operations MIB
- RFC 2932 IPv4 Multicast Routing MIB
- RFC 2933 IGMP MIB
- RFC 2934 PIM MIB
- RFC 3019 IPv6 MIB: MLD Protocol
- RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
- RFC 3412 Message Processing and Dispatching for the Simple Network Management Protocol
- RFC 3416 v2 of the Protocol Operations for the Simple Network Management Protocol
- RFC 3635 Definitions of Managed Objects for the Ethernet-like Interface Types
- RFC 3636 Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units
- RFC 3810 Multicast Listener Discovery v2 for IPv6
- RFC 3811 Definitions of Textual Conventions for Multiprotocol Label Switching Management
- RFC 3812 Multiprotocol Label Switching Traffic Engineering Management Information Base
- RFC 3813 Multiprotocol Label Switching Label Switching Router Management Information Base
- RFC 3815 Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol
- RFC 4022 Management Information Base for the Transmission Control Protocol 4087 IP Tunnel MIB
- RFC 4113 Management Information Base for the User Datagram Protocol
- RFC 4624 Multicast Source Discovery Protocol MIB
- RFC 4893 BGP Support for Four-octet AS Number Space
- RFC 6329 IS-IS Extensions supporting IEEE 802.1aq Shortest Path Bridging
- draft-unbehagen-spb-ip-ipvpn-00 - IP/IPVPN Services with IEEE 802.1aq SPB networks

Weights & Dimensions

- Ethernet Routing Switch 8010 Chassis - 14RU
Height: 22.9 in. (58.2 cm)
Width: 17.5 in. (44.5 cm)
Depth: 19.9 in. (50.5 cm)
Weight: up to 225 lb (102 kg)
Cooling System:
- Fan Trays: 2 per Chassis
- Fans: 15 per Fan Tray
- Thermal Sensors: 1 per Fan Tray
- Ethernet Routing Switch 8010co Chassis - 20RU
Height: 35.9 in. (88.9 cm)
Width: 17.5 in. (44.5 cm)
Depth: 23.7 in. (60.2 cm)
Weight: up to 315 lb (143 kg)
Cooling System: 8010co features front-to-back cooling and the maximum airflow specification for the 8010co Chassis is 330 linear ft/min

The 8010co Chassis complies with Network Equipment Building Standard (NEBS) Level 3 as specified in SR3580. Typically the 8010co has superior physical and environment specifications; please reference to the product documentation for full details

- Ethernet Routing Switch 8006 Chassis - 10RU
Height: 15.8 in. (40.1 cm)
Width: 17.5 in. (44.5 cm)
Depth: 19.9 in. (50.5 cm)
Weight: up to 170 lb (77 kg)
Cooling System:
- Fan Trays: 1 per Chassis
- Fans: 20 per Fan Tray
- Thermal Sensors: 1 per Fan Tray
- Ethernet Routing Switch 8003R Chassis - 7RU
Height: 12.25 in. (31.1 cm)
Width: 17.5 in. (44.5 cm)
Depth: 21.0 in. (53.5 cm)
Weight: up to 76 lb (34.5 kg)
Cooling System:
- Fan Trays: 1 per Chassis
- Fans: 3 per Fan Tray

Environmental Specifications

- Operating temperature: 0°C to 40°C (32°F to 104°F)
- Storage temperature: -25°C to 70°C (-13°F to 158°F)

- Operating humidity: 85% maximum relative humidity, non-condensing
- Storage humidity: 95% maximum relative humidity, non-condensing
- Operating altitude: 3024 m (10,000 ft) maximum
- Storage altitude: 3024 m (10,000 ft) maximum
- Free fall/drop: ISO 4180-s, NISTA 1A
- Vibration: IEC 68-2-6/34
- Shock/bump: IEC 68-2-27-29

Safety Agency Approvals

- Global basis for certification: IEC 60950 current edition with all CB member deviations
- US: UL60950
- Canada: CSA 22.2 No. 60950
- Europe: EN60950 (CE Marking)
- Australia/New Zealand: AS/NZS 3260
- Mexico: NOM-019-SCFI-1998

Electromagnetic Emissions

- Global basis for certification: CISPR 22-1997 Class A

- US: FCC CFR47 Part 15, Subpart B, Class A
- Canada: ICES-003, Issue-2, Class A
- Europe: EN 55022-1998 Class A; EN 61000-3-2/A14

Electromagnetic Immunity

- Global basis for certification: CISPR 24:1997
- Europe: EN 55024:1998

Warranty

- 12 months on principle components

MTBF Values

- Chassis: 653,732 to 2,043,676 hours
- Power Supplies: 108,803 to 282,805 hours
- Switch Fabric/CPU Module: 213,454 hours
- Interface Modules: 350,814 to 388,833 hours

Country of Origin

- Malaysia (with the exception of Power Supplies which are manufactured in PR China)

Ordering Information (principle components)

Component	Part Code	Description
Chassis	DS1402001-E5	8010 10-Slot Chassis, complete with Fan Trays
	DS1402004-E5	8010co 10-Slot Central Office Chassis, complete with Fan Trays
	DS1402002-E5	8006 Series 6-Slot Chassis, complete with Fan Trays
	DS1402011-E5	8003R Series 3-Slot Chassis, complete with Fan Tray
System Modules	DS1404120-E5	8895SF Switch Fabric/CPU Module
Power Supplies	DS1405012-E5	8005AC 100-240V AC Power Supply, up to 1462W
	DS1405018-E5	8005AC 100-240V AC Dual Input Power Supply, up to 1462W
	DS1405011-E5	8005DC DC Power Supply, 1462W
	DS1405017-E5	8005DC DC Dual Input Power Supply, 1462W
Interface Modules	DS1404121-E6	8812XL 12-port 10GbE SFP+ Ethernet Interface Module
	DS1404123-E6	8834XG 34-port Combo 10/100/100 & 1GbE SFP & 10GbE XFP Ethernet Interface Module
	DS1404122-E6	8848GB 48-port 1GbE SFP Ethernet Interface Module
	DS1404124-E6	8848GT 48-port 10/100/1000 Ethernet Interface Module
Software	DS1410003-7.2	ERS 8800 Base Software License
	DS1410021	ERS 8800 Advanced Software License, activating additional features: BGP4, IPv6 Routing, BFD, MSDP, & PCAP
	DS1410026	ERS 8800 Premier Software License, activating Advanced features, plus Fabric Connect (including Multicast), IP VPN-Lite, MP-BGP, MPLS IP VPN, & VRF-Lite (including Multicast)

About Avaya

Avaya is a global provider of business collaboration and communications solutions, providing unified communications, contact centers, networking and related services to companies of all sizes around the world. For more information please visit www.avaya.com.