

The VSP 7000's capabilities have been enhanced with the 10.3 feature releases, including:

- Enhancing the Avaya VENA Fabric Connect technology - our end-to-end network virtualization capability - enabled for concurrent interoperability with the Switch Cluster technology
- Introducing the new VSP 7024XT 24-port 10GBASE-T Switch
- Introducing the new 7002QQ 2-port 40 Gigabit QSFP+ Media Dependent Adapter
- Introducing our Software-Defined SAN technology for converged storage support
- Re-structured Software Licensing delivers all features in the default Base License for enhanced ownership costs

For more information about the various Avaya VENA capabilities, please visit www.avaya.com/vena.

Avaya Virtual Services Platform 7000 Series

Future-ready Ethernet switching platform specifically architected for tomorrow's high-performance Data Center

The Avaya Virtual Services Platform 7000 Series, with its unique 'Distributed Top-of-Rack' capability, is the Data Center Top-of-Rack solution that delivers the industry's best application performance; featuring multi-hop low-latency, supporting advanced Fabric-based services, and enabling a one-touch, edge-only provisioning model. Leveraging this application performance advantage, businesses can reduce time-to-decision cycles, and reduce both deployment and operational costs.

The Avaya Virtual Services Platform 7000 Series is an energy and space-efficient platform built around a next-generation chipset, delivering wire-speed 1/10Gbps Ethernet for today's connectivity requirements. It is also uniquely future-ready with the embedded flexibility to seamlessly support 40 and 100Gbps Ethernet, Storage convergence, and has been built to support network-wide fabric-based virtualized services and lossless environments.

Highlights

- Built around an innovative fifth generation flexible high-performance ASIC chipset with optimized packet buffering and enlarged scaling capabilities
- Versatile cooling options conforming to established data center hot/cold aisle arrangements
- Hot-swappable power supplies, fans, and pluggable expansion
- Wire-speed hardware guaranteeing high throughput and low latency

- Dedicated high-speed ports empower the flexible, low-latency Distributed Top-of-Rack capacity that seamlessly connects multiple units
- Future-ready support for sophisticated virtualization technologies and high-speed interfaces

Ready for today

The Virtual Services Platform 7000 Series (VSP 7000) is ideally suited to deliver today's pressing need for flexible, high-speed Ethernet connectivity in the high-performance data center Top-of-Rack (ToR) role. Additionally it provides a cost-effective 10 Gigabit Ethernet fan-out capability for existing core switch deployments, saving valuable ports and slots. It also has application in the campus distribution layer, delivering flexible connectivity and consolidation options. Featuring a hardened physical architecture of dual, hot-swappable AC or DC power supplies and fan trays, the VSP 7000 is an important addition to a network manager's toolkit for creating always-on high-performance solutions.



VSP 7024XLS Ethernet Switch with MDA Slot



VSP 7024XT Ethernet Switch with MDA Slot



VSP 7000 Series rear view - note the four Fabric Interconnect ports (located between the removable power supplies and fans)

The VSP 7000 Series includes two variants. The 7024XLS, the first model introduced, features 24 fixed SFP+ sockets, supporting a wide variety of both 1Gbps and 10Gbps Ethernet pluggable transceivers, enabling short- and long-haul copper and fiber connections. The new 7024XT model features 24-ports of fixed 10GBASE-T support, presented as RJ45 connectors.

In addition, the VSP 7000 Series products also feature a versatile Media Dependent Adapter slot that supports a range of high-speed expansion options, such as additional 10Gbps ports (SFP+ sockets or 10GBASE-T via RJ45), 40Gbps and 100Gbps¹ Ethernet connectivity.

An innovative design helps ensure that the appropriate portion of the



7002QQ Media Dependent Adapter featuring 2 x 40GBASE-QSFP+ sockets



7008XLS Media Dependent Adapter featuring 8 x 10GBASE-SFP+ sockets



7008XT Media Dependent Adapter featuring 8 x 10GBASE-T ports

switching fabric's powerful 1,280Gbps performance is dedicated to supporting MDA-based connections, while still providing for wire-speed throughput for all front panel and high-capacity interconnects.

The design of the VSP 7000 is sympathetic to the evolving requirements for precisely planned environmental implementations and it offers both front-to-back and back-to-front options for the field-replaceable fans tray and power supplies; this gives the flexibility to conform to a variety of hot-aisle/cold-aisle design requirements.

Future-ready for tomorrow

In addition to the versatility afforded by the optional MDA, the VSP 7000 includes a host of strategic innovations that help ensure it will continue to deliver cutting-edge performance and services throughout an extended life cycle. Among the future-ready features is the advanced fifth generation application-specific integrated circuit (ASIC) chipset that has been optimized for network operations, including:

- Smart packet buffering to deliver lower latencies and efficiently reallocate memory to those ports suffering congestion
- Larger table capacities provide for improved scaling
- Embedded services in support of applications such as IP Flow Information Export (IPFIX)
- Native support for the particular computational functionality crucial in high-performance data center and metro deployment scenarios
- Native capability for multiple advanced technologies, with possible support for the likes of DCB, TRILL, MPLS, VPLS and the associated OA&M; specific support being driven by market demand.

In addition to a class-leading switching performance of over 1.2Tbps, the VSP 7000 has been equipped from the outset with superior CPU and memory resources so that it has the capabilities

to execute the sophisticated algorithms necessary for a successful fabric-based infrastructure.

Deployment options and benefits:

- Scenarios:
 - » Top-of-Rack Data Center supporting single/dual-attach
 - » Fibre aggregation Distribution Switch or Data Center 10 Gigabit fan-out
 - » Ethernet Storage transport leveraging low-cost ATA-over-Ethernet
 - » Fibre Channel-over-Ethernet integration with a Software-Defined SAN capability
- Scope:
 - » Flexible DC deployment scenarios, either traditional or Fabric-based
 - » Transitional between 1G and full 10G plant, seeking a strategic solution
- Benefits:
 - » 1/10GbE support
 - » Future-ready for key enabling technologies
- Positioning:
 - » Seamlessly transition the highly virtualized server environment from multiple 1GbE connections to one or two 10GbE - increase capacity, reduce latency, reduce complexity
 - » Make a strategic investment in a future-ready platform that matches the mainstream data center evolution requirements and timeline - 40/100GbE, Storage convergence, DCB
 - » Class-leading performance package - lowest latency, highest switching, most advanced chipset, and unique Fabric Interconnect ports that are built-in and support flexible Distributed Top-of-Rack deployment scenarios

¹ Indicates a roadmap feature that is forecast to be delivered in a future software release.

Avaya VENA Distributed Top-of-Rack

The Avaya VENA Stackable Chassis feature is a common thread that runs through our fixed-format Switching products. Built using Avaya's pioneering 'Flexible Advanced Stacking Technology' (FAST), this capability delivers genuine chassis-like levels of resiliency and performance at a fixed format price. This is one of the differentiators that make Avaya genuinely unique to competitive offerings. The experience gained developing successive generations of this technology across multiple product lines has enabled Avaya to now deliver the 'Distributed Top-of-Rack' capability with the VSP 7000 Series. This capability creates a virtual backplane optimizing local traffic flows by always forwarding packets along the shortest path between source and destination, using dedicated high-speed interconnections. This is particularly beneficial in a data center environment creating a flatter, latency-free network that encompasses multiple racks of highly virtualized servers. Available originally as a scale-out implementation of FAST, Distributed Top-of-Rack was extended with an additional mode - Fabric-mode - and this increases deployment flexibility and operational agility.

Stack-mode: the same proven capability, now faster

High-capacity virtual backplane - it's been a given that performance comes as a natural function of design and price. However, with Stack-mode Distributed Top-of-Rack, Avaya has been able to combine non-blocking internal Switching Fabrics with a high-speed interconnectivity architecture to deliver a truly high performance and cost-effective solution. The Stack-mode implementation is not bound by the limitations and constraints facing rivals (such as token sharing/passing systems, or basic cascading), and has been specifically designed to scale proportionally as new member switches are added; as more ports



are added and the requirement for more bandwidth grows. A shortest-path traffic forwarding capability is at the heart of our resilient architecture, ensuring that the shortest, most optimal forwarding path is selectively chosen for each unique data flow. There is none of the unwieldy logical ring or token technology that lesser offerings use, but a star-based distributed forwarding topology that allows traffic to flow either 'upstream' or 'downstream' simultaneously from every Switch in the system, optimizing both performance, resiliency, and resource utilization. The VSP 7000 features our most advanced implementation to date, delivering a familiar, proven capability that can be used as a scalable building block for an evolving, high-performance Data Center environment.

In-service maintenance and restoration

Virtual hot-swap capability - this crucial serviceability and operability feature helps ensure that any unit failure can be quickly and easily rectified. It's a hot-swap capability pioneered in modular switches, and also made available on VSP 7000 when deployed in Stack-mode. Enabling immediate like-for-like unit replacement without impacting other functionality and traffic, and without complex engineering intervention, empowers operators to deploy our solutions just as they would a chassis. When a failure occurs the neighboring switches will automatically wrap their fabric connections to help ensure that adjacent racks in the data center are not impacted. The failed unit is simply disconnected from the virtual

backplane and a like unit - without any pre-staging of software or configuration - can be inserted, cabled, and powered-up. The Automatic Unit Replacement (AUR) process self-manages any necessary downloads to the new switch and then brings it online; all of this without the need for an engineer to configure or manage the process.

Distributed power and forwarding

No single point-of-failure - with a chassis solution this has meant N+1 power supplies and even redundant Switching Fabrics. With Distributed Top-of-Rack implemented in Stack-mode it is much the same but without the cost penalty: each unit has an independent Switching Fabric, and each 'Switching Module' has an independent power supply, which means that there is no one single point-of-failure. The VSP 7000's dual hot-swappable power options further enhances the overall resiliency of the solution. Any individual element failure is equivalent to the failure of a single module within a chassis. Frame forwarding decisions are distributed across the Fabric; when a frame forwarding decision results in the need to forward the frame to another Switch, the intelligent shortest path algorithm determines whether the frame will traverse the virtual backplane.

Automatic software & configuration control, and centralized management

Simplified implementation and management - with a chassis solution it is simply a case of adding a new module, adding configuration,



and connecting devices. With the VSP 7000, it is much the same: simply cable-in a new member, extend the appropriate configuration – all units are managed as a single network entity. The Automatic Unit Replacement (AUR) feature is the process that delivers the Agent Image software, the configuration file, and the Diagnostic Image software to any additional or replacement Switch; it self-manages any necessary downloads to the new Switches and then brings it online. All units operate, appear, and are managed as a single network entity, with unique IP and MAC Addresses.

Management functions are performed by a 'Master Unit' (MU), which is selected as part of initializing the virtual backplane. Once normal operations have been established, if the MU (e.g. Unit 1) fails or is rebooted, the next Switch (e.g. Unit 2) will take over as Temporary MU (TMU) and remain as such until either it or the entire system is reset. If Unit 2, as the TMU, fails or is reset, then the TMU status will transfer to the next downstream unit. The important point is that there is always a process for electing the master and a backup to help ensure continuous operation of the combined system.

Fabric-mode: enhanced functionality for next-generation, scale-out virtual networking

The VSP 7000 is an integral component of our Fabric Connect strategy for end-to-end Fabric-enabled networking, and the introduction of Shortest Path

Bridging (SPB) enables us to leverage this technology to further innovate by developing the unique Distributed Top-of-Rack capability. With Fabric-mode, multiple VSP 7000s can be very flexibly meshed to uniquely deliver multi-hop, low-latency for deployment scenarios that call for massive scale-out.

Avaya has architected the Fabric-mode Distributed ToR solution to ultimately scale up to many hundreds of Switches within a single domain, with the current implementation supporting up to 500 Switches networked as a single contiguous Fabric. This configuration delivers scale of up to 16,000 wire-speed 10 Gigabit Ethernet ports - or up to 12,000 ports of 10 Gigabit and 1,000 ports of 40 Gigabit - that are directly supported by a virtual backplane with a capacity of 280Tbps. The agile building block approach is extremely flexible, with virtually no hard-and-fast topology constraints; blocks can be small or large, and individual Switches easily interconnected with extended-reach copper or fiber cabling.

The Fabric-mode of operation differs from the existing Stack-mode option in a number of key areas:

- The Shortest Path Bridging standard is used as the internetworking protocol technology (as opposed to the Avaya FAST protocol) to control the forwarding of traffic between Switches
- The upper scale of a single 'domain' is determined by the number of SPB Nodes supported within software – the current supported

number is 500 – as opposed to the 'hard-coded' FAST limit of 8 (per system)

- The virtual backplane capacity increases from the Stack-mode limit of 5.12Tbps, and is up to 280Tbps (based on 500 Switches)
- Like-for-like Switch replacement, i.e. how replacement units are introduced and how image and configuration files are deployed, is a manual process in Fabric-mode, as opposed to the automated functionality delivered by Stack-mode Auto Unit Replacement
- The rules governing how Fabric Interconnect ports are cabled changes from the strict requirements of Stack-mode to 'free' with Fabric-mode, promoting very flexible topologies
- Switches are individually managed in Fabric-mode as opposed to the group management approach with Stack-mode
- And finally, the configuration effort required to deploy both modes is relatively similar with the very low configuration burden being a feature of both

As can be seen, the expansion of the Distributed Top-of-Rack modes of operation provides powerful additional capabilities and deployment flexibility. Some scenarios may call for Fabric-mode's higher interconnection scale or more flexible cabling approach while others might take advantage of Stack-mode's additional active-active options or streamlined management.

Avaya VENA Switch Cluster

The Avaya VENA Switch Cluster technology is built using the Split Multi-Link Trunking technology that is unique to our products, yet is fully interoperable with third party Switches, Servers, Appliances, and Routers. What this delivers is a series of benefits that provide real value; while it may be possible to simulate certain individual elements, there is no competitive offering that can rival the

combined capabilities, particularly in terms of simplicity and efficiency.

Delivering end-to-end application availability

Interoperable solution that extends beyond Switches to Servers – this means that the high-availability is not limited to only the switching network (the Switches themselves and their direct links), but to the total network; importantly also extended to attached Servers, Appliances, and WAN Routers, etc. All competing offers are based on interaction within the Switch domain, and crucially do not extend to the application hosts themselves. Most rival offerings are based around variations of the Spanning Tree Protocol (STP); however this is limited to the actual switches and is not supported by other devices (servers, etc). By excluding servers from the active resiliency technology, these solutions cannot extend availability to the applications themselves. The Avaya VENA Switch Cluster technology is independent of STP and extends to support any device that utilizes Link Aggregation, a technology that is both basic and ubiquitous. Devices that attach to the Switch Cluster create a virtual connection using multiple physical links, this provides resiliency together with additional capacity.

Solutions that scale enterprise-wide

Cost-effective solutions for every network size, not limited to top-of-the-range product – business processes that demand high-availability are typically also deployed company-wide; so it is imperative to deliver consistent levels of resiliency across the entire network. Avaya is uniquely positioned to offer the same capability, using the same technology, in a broad range of platforms that scale in both price and performance matching various requirements across the network. This enables the consistent delivery of a network that is itself constantly delivering end-to-end application availability, regardless of location or size.

Simplified solutions

A capability that is simple to deploy and one that does not require complex and/or expensive products – the various capabilities and benefits that Switch Cluster technology provides are delivered without complexity. There is no need for expensive hardware or software, or for complex configuration or on-going maintenance. This helps ensure that the business benefits are not ‘purchased at any price’; indeed Switch Cluster technology uniquely combines simplicity with cost-effectiveness. While some of the individual capabilities can be simulated using other techniques and/or a myriad of additional products, that approach can only add cost and complexity.

As an example, Switch Cluster technology delivers user pre-session load-sharing across all uplinks from the access switch to the core; this capability is automatically enabled without the need for any additional hardware, software, or configuration. To attempt to replicate this level of capability in a STP-based network, it would need to have multiple VLANs with MSTP/RSTP configured, a Layer 3-enabled Switch with routing and ECMP configured, and even then this solution would be limited to pre-VLAN load-sharing only (not pre-session).

Sub-second failover & recovery

Delivering the necessary availability and also facilitating in-service maintenance and optimized performance – Switch Clustering is probably best known for delivering sub-second failover and recovery. While this remains extremely important, and never more important, it is not necessarily a feature that remains unique. Enhancements to STP – namely rapid reconfiguration – can be aggressively configured to deliver similar levels of failover performance. However all flavors of STP remain tied to the concept of detecting and acting upon changes to the network topology. This makes a network extremely sensitive to the reliability

and availability of particular devices (root bridges, etc). Avaya’s Switch Cluster technology is built around the concept of mirrored devices and virtualized capabilities, so that an entire switch can be removed (through failure or for emergence or routine maintenance) without any loss of overall application availability.

What Avaya continues to deliver in this area is a degree of network recovery which also facilitates in-service maintenance. The deterministic nature of Switch Cluster technology empowers network operators to compartmentalize the network, making essential services even more resilient, and allowing for individual failures to be repaired in real time, without service restoration work impacting on collateral components or applications.

Virtualization and network fabric infrastructure

Most enterprises are looking at ways to reduce cost and improve time-to-service, and most are looking for ways to improve the operational efficiency of the Data Center. Wouldn’t it be nice if you could deploy a new application across multiple data centers in an instant? Wouldn’t it be nice if you could give IT the ability to simply say “yes” when there is a new application or service that needs to be deployed on the network? Server virtualization within the data center is now taken for granted, with some declaring that ‘Cloud Computing’ will become a reality for most enterprises, and that applications, information, and compute resources will become simple commodities. Experience has proved one thing; the Data Center of the future cannot be built on the technology of the past. General-purpose products, outmoded techniques, and legacy designs cannot be re-packaged as ‘data center-ready’. The industry will take the best and leave the rest. Ethernet is readily available, cost-effective, extensible, and – as the 40 and 100 Gigabit developments prove –



scalable, however many existing deployment methodologies are no longer an option.

The Avaya Virtual Enterprise Network Architecture is an overarching strategy designed to improve network uptime by delivering the infrastructure that creates the private cloud and virtually eliminates user-error network outages. We are reducing time-to-service with simple one-touch provisioning and we are improving Data Center efficiency with a tight integration between applications and network virtualization. We are reliably connecting users and content, with independent testing commissioned by Avaya reinforcing the fact that we consistently deliver some of the best total cost of ownership in the industry. With an Avaya data network, you will get more value for every dollar spent by IT.

Avaya VENA Fabric Connect

The Avaya VENA Fabric Connect technology, based on an extended implementation of the Shortest Path Bridging (SPB) standard, offers the ability to create a simplified network that can dynamically virtualize elements to empower efficient provisioning and utilization of resources, thus reducing the strain on the network and personnel. Combined with the Avaya Virtualization Provisioning Service tool, the Avaya VENA Fabric Connect technology can, for example, fully synchronize the provisioning of the networking requirements of workload mobility within and between data centers, dynamically and seamlessly moving

or extending virtualized computing resources, without the provisioning complexity associated with rival solutions. Based on SPB, an IEEE 802.1aq standard augmented with Avaya enhancements that deliver specific enterprise optimization, Fabric Connect offers a robust and resilient alternative to today's existing offerings and it delivers innovative services and solutions while maintaining Ethernet's key value propositions of simplicity and cost-effectiveness. Fabric Connect delivers new capabilities in the crucial areas of simplicity, scalability, performance, reliability, and service orchestration and abstraction.

Creating a fault-tolerant, powerful, and self-aware end-to-end fabric, this transparent network features a design where service provisioning occurs only at the perimeter. The advantage is immediate and pronounced; administrative effort is reduced, errors are avoided, and time-to-service is vastly enhanced. The beauty of the underlying SPB technology is that it masks devices, links, and protocols and delivering what is logically an extended Ethernet LAN that provides connectivity for multiple end-points. That's the simple concept and SPB achieves this in an interesting and quite unique way. It leverages a dynamic link-state routing protocol called Intermediate System-to-Intermediate System (IS-IS) and extends it to share topology, reachability, and device information between every node in the SPB domain. With nodes holding their own self-determined view of the network, including the optimal path to any destination, a fully distributed and dynamically maintained solution is created.

Avaya Virtualization Provisioning Service

Avaya Virtualization Provisioning Service is a virtualization management solution that delivers automation, visibility & reporting that spans the network infrastructure, servers, storage and applications, across both physical and virtual environments. Today's Data Center networks are inefficient in dealing with server virtualization. From the time a Virtual Machine (VM) is created to the time it is activated, moved around or deactivated, the network has no visibility into the virtual machine lifecycle. There are also few tools, if any, when it comes to troubleshooting and managing VMs in the network. Due to the ability of VMs to dynamically move from server to server, provisioning the network for VM security and application performance has proven to be a very serious networking challenge.

An important milestone in the evolution to a virtualized Data Center is making the network very "efficient" when it comes to managing, troubleshooting, provisioning and securing virtual machines in the network. This includes bringing network level insight and visibility to the virtual machine lifecycle, applying the appropriate network and port level configurations at an individual VM level, dynamically tracking VMs as they move across the data center and enforcing the network attributes of the VMs wherever they migrate in the data center. Avaya's Virtualization Provisioning Service (VPS) is a software application service that acts as glue between VMware's vCenter and Avaya's Configuration and Orchestration Manager. Avaya Virtualization Provisioning Service provides a relay mechanism to bridge the gap of complete end-to-end provisioning of servers and network devices in a fully virtualized Data Center environment. It learns dynamic virtualized server topologies and updates Avaya devices to react to changes in server topologies. It

provisions connectivity services (VLANs) on Switch ports based on the actual network connectivity and provisions QoS filter, ACL, SPB I-SID & port profiles (templates) – based on preconfigured rules – to Switch ports.

Provisioning changes can be applied automatically based on a set of predefined rules which are checked and applied to network ports dynamically if the rule applies. They can also be applied manually where there is an alert to the network administrator that there is a change happening within the server environment and then a guided workflow is triggered that would allow the administrator to apply the manual network configuration change. These changes to the network, which can be done in real-time, are critical in helping to ensure the applications function as expected and that moving the VM doesn't negatively impact the end-user experience for that particular application,

An area of differentiation for Avaya Virtualization Provisioning Service is its integration with the Avaya Identity Engines portfolio to deliver intelligent rule-based access control for individual VMs. This gives network administrators the ability to prevent individuals from moving certain VMs in the middle of the day increases protection of specific networks so that only approved VM's can be connected to them.

Avaya Virtualization Provisioning Service will also provide a wealth of reporting options so that network operators have a clear view of the VM lifecycle and activity (activations, deactivations, changes), it will provide details on what network changes were completed based on user, device, time, type of access etc. Network operators can also customize the alerts that they receive based on the device type, port groups or even server type.

This comprehensive solution truly brings the virtualized applications together with the virtualized network and helps ensure that the network is able to constantly adapt to changes in the computing environment. VM

mobility is then transparent to the end user utilizing those applications.

Management

The VSP 7000 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 evolving Unified Management framework for comprehensive, centralized, and multi-faceted network management. It 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 & Orchestration Management; Visualization, Performance and Fault Management; and IP Flow Manager.

Provision wizards and other labor-saving tools help ensure 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 mandatory 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.

Software-Defined Networking

Enhanced agility, flexibility, time-to-service, and automation are the main aspirational goals of Software-Defined Networking (SDN). Avaya's pragmatic approach is focused on delivering solutions for real-world business challenges.

Software-Defined Data Center (SDDC)

Avaya's Software-Defined Data Center framework is designed to deliver productivity, agility and time-to-service enhancements to businesses operating highly virtualized Data Centers. The SDDC framework aims to break down traditional Data Center silos that require weeks or months to turn up an application and replace it with a simple five-step process that takes minutes.

Avaya's framework includes an orchestration process that combines, customizes and commissions compute, storage and network components. Use of the OpenStack cloud computing platform will allow Data Center administrations to deploy virtual machines, assign storage and configure networks through a single graphical user interface. Fabric Connect further enhances the OpenStack environment by removing restrictions in traditional Ethernet VLAN/Spanning Tree-based networks to enable a more dynamic, flexible and scalable network services model than exists today.

The Avaya SDDC framework is based on the following components:

- Fabric Connect technology as the virtual backbone to interconnect resource pools within and between data centers with increased flexibility and scale
- An Avaya-developed OpenStack Horizon-based Management Platform, delivering orchestration for compute (Nova), storage (Cinder/Swift) and network via Avaya Fabric Connect (Neutron)
- Open APIs into the Fabric Connect architecture for ease of integration, customization and interoperability with other Software-Defined Networking architectures

Ultimately, Avaya's SDDC framework can ultimately provide the following benefits to businesses:

- **Reduced Time-to-Service:** cloud services enabled in minutes through a few simple steps
- **Simplified Virtual Machine Mobility:** simple end-point provisioning enabling virtual machine mobility within and between geographically dispersed data centers
- **Multi-Vendor Orchestration:** coordinated allocation of data center resources via a single interface to streamline the deployment of applications
- **Scale-out Connectivity:** services scale to more than 16 million unique services, up from the limitation of 4000 in traditional Ethernet networks
- **More Secure Multi-Tenancy:** leveraging network, compute and storage layer abstraction and isolation
- **Improved Network Flexibility:** overcomes current VLAN challenges to deliver a load-balanced, loop-free network where any logical topology can be built, independent of the physical layout

The VSP 7000 Series is fully integrated into Avaya's SDDC framework, with device configuration orchestrated via the Avaya SDN Controller plug-in component of the OpenStack-based Horizon dashboard.

Software-Defined Storage Area Networking (SDSAN)

Avaya's Software-Defined Storage Area Network provides a compelling solution for businesses that wish to combine existing Fibre Channel/Fibre Channel over Ethernet (FCoE) investments and expertise with the efficiencies of a next-generation consolidated Ethernet network. De-coupling and centralizing the FCoE control plane to a virtualized FCoE Network Controller (FNC) significantly simplifies the infrastructure requirements, leveraging the many advantages of Fibre Channel, FCoE, and Ethernet.

The FNC is a software engine that emulates the functions of a traditional Fibre Channel Forwarder (FCF), delivering seamless compatibility but avoiding all of the complexity and costs associated with traditional hardware-based solutions. Integrated with the VSP 7000 Series, the FNC creates a truly software-defined solution; FCoE control messages are interpreted by the FNC and the necessary traffic control policies are dynamically written to the VSP 7000. In this model, a full-feature FCoE-based Storage Area Network (SAN) is created without any additional Fibre Channel or FCoE hardware; it's a purely software-based and software-defined solution.

Another advantage of the FNC is that, being a software engine, it can be flexibly deployed on existing Virtual Machine (VM) platforms; indeed, this virtualized model provides for enhanced level of availability and resiliency relative to hardware-based offerings.

Avaya is introducing the SDSAN technology with a level of scalability suitable for small- to medium-sized network deployments. Initially, the technology will support FCoE connectivity consolidated to a single Switch/Switch pair, with an associated FNC; essentially up to 32 single/dual-attached physical nodes. This relatively conservative approach ensures that loss-sensitive storage is not compromised by the 'best-efforts' nature of general purpose Ethernet networks. As the Data Center Bridging standards evolve, we expect to be able to expand the SDSAN capability to support multi-Switch/multi-hop deployments.

Lifetime warranty

Avaya includes Industry-leading warranty services for the portfolio of Stackable Chassis products, including the VSP 7000. We provide complimentary next-business-day shipment of failed units for the full life of the product; next-business-day

shipping to replace failed hardware worldwide. Avaya also offers complimentary basic technical support: Level 1 the supported lifecycle of the product and up to Level 3 for the first 90 days after purchase; this includes support for the shipped software version, with an optional Software Release Service available to provide access to new feature releases. As per industry norm for hardware, 'Lifetime' is defined as the production lifecycle phase, plus 5 years post-discontinuation.

Summary

The Avaya Virtual Services Platform 7000 is purpose-built to support the dynamic data center and high-density 10 Gigabit Ethernet Top-of-Rack deployments of today. It alleviates infrastructure complexity and reduces power consumption with a truly scalable and strategic architecture; it is designed to be the high-performance top-of-rack platform for the next decade.

Supporting mission-critical applications requires 24/365 always-on infrastructure, and the VSP 7000 delivers against this challenge. It is a highly strategic product that is fit-for-purpose for today's connectivity requirements and future-ready for the evolving and emerging application-driven needs of tomorrow.

The VSP 7000 brings unique differentiation to the ToR role: with a flexible, non-blocking architecture, including wire-speed Server access connections and Fabric Interconnect ports. The VSP 7000 is purpose-built to support today's dynamic Data Center operations and high-density, low-latency 10 Gigabit Ethernet Top-of-Rack deployments. It alleviates infrastructure complexity and reduces power consumption with a truly scalable and strategic architecture; it is designed to deliver a high-performance Distributed Top-of-Rack solution that fully optimizes next-generation application virtualization investments.

Ordering Information

Part Code	Product Description
AL700001B-E6	Avaya Virtual Services Platform 7024XLS 24-port 1/10 Gigabit Ethernet SFP+ Switch (Back-to-Front Cooling)
AL700001F-E6	Avaya Virtual Services Platform 7024XLS 24-port 1/10 Gigabit Ethernet SFP+ Switch (Front-to-Back Cooling)
AL7000S2B-E6	Avaya Virtual Services Platform 7024XT 24-port 10 Gigabit Ethernet Switch (Back-to-Front Cooling)
AL7000S2F-E6	Avaya Virtual Services Platform 7024XT 24-port 10 Gigabit Ethernet Switch (Front-to-Back Cooling)
AL7000MV2-E6	7002QQ 2-port 40GBASE-QSFP+ Media Dependent Adapter
AL7000MS1-E6	7008XLS 8-port 10GBASE-SFP+ Media Dependent Adapter
AL7000MT1-E6	7008XT 8-port 10GBASE-T Media Dependent Adapter
AL7000?0F-E6	Avaya Virtual Services Platform 7000 AC Power Supply (Front-to-Back Cooling)
AL7000?0B-E6	Avaya Virtual Services Platform 7000 AC Power Supply (Back-to-Front Cooling)
AL7000A1F-E6	Avaya Virtual Services Platform 7000 DC Power Supply (Front-to-Back Cooling)
AL7000A1B-E6	Avaya Virtual Services Platform 7000 DC Power Supply (Back-to-Front Cooling)
AL7000FTB-E6	Spare Fan for Avaya Virtual Services Platform 7000 (Front-to-Back Cooling)
AL7000BTF-E6	Spare Fan for Avaya Virtual Services Platform 7000 (Back-to-Front Cooling)
AL7018001-E6	Fabric Interconnect Cable (Copper) - 0.6m
AL7018002-E6	Fabric Interconnect Cable (Copper) - 1.0m
AL7018003-E6	Fabric Interconnect Cable (Copper) - 3.0m
AL7018004-E6	Fabric Interconnect Cable (Copper) - 5.0m
AL7018005-E6	Fabric Interconnect Cable (Fiber) - 10.0m (with fixed Transceiver ends)
AL7018006-E6	Fabric Interconnect Fiber Transceiver for use with 50m & 100m Fabric Interconnect Cables (Fiber) - 2 x Transceivers required per Cable
AL7018007-E6	Fabric Interconnect Cable (Fiber) - 50m - Requires 2 x Fabric Interconnect Fiber Transceivers (AL7018006-E6)
AL7018008-E6	Fabric Interconnect Cable (Fiber) - 100m - Requires 2 x Fabric Interconnect Fiber Transceivers (AL7018006-E6)
AL7011001-E6	Avaya Virtual Services Platform 7000 4 Post Server Rack Mount Kit
FC7000D32	FCoE Controller License for up to 32 FCoE Devices
Where applicable the seventh character (?) of the Order Code is replaced to indicate the required product nationalization:	
A	No Power Cord option
B	Includes European "Schuko" Power Cord option, common in Austria, Belgium, Finland, France, Germany, Netherlands, Norway and Sweden
C	Includes Power Cord used in UK and Ireland
D	Includes Power Cord used in Japan
E	Includes Power Cord used in North America
F	Includes Power Cord used in Australia, New Zealand and People's Republic of China

Specifications

General and Performance

- Physical Connectivity:
 - 7024XLS
 - 24 x 10GBASE-SFP+ Ports (front)
 - 1 x Media Dependent Adapter Slot (front)
 - 4 x High-Speed Fabric Interconnect Ports (rear)
 - 7024XT
 - 24 x 10GBASE-T RJ45 Ports² (front)
 - 1 x Media Dependent Adapter Slot (front)
 - 4 x High-Speed Fabric Interconnect Ports (rear)
- Switch Fabric Architecture: 1,280Gbps Full-Duplex
- Frame forwarding rate: 960Mpps per Switch

- Typical Latency: ~500nsec
- Typical Jitter: 12 - 14Qsec
- Frame length: 64 - 1518 Bytes (Untagged), 64 - 1522 bytes (802.1Q Tagged)
- Jumbo Frame support: up to 9,216 Bytes (802.1Q Tagged)
- Multi-Link Trunks: up to 64 Groups, with 8 Links per Group
- VLANs: up to 1,024 (up to 4,094 VLAN IDs)
- Multiple Spanning Tree Groups: 8
- MAC Address: 131,000
- DHCP Relay Entries: up to 256
- ARP Entries: up to 4,096
- IP Interfaces: up to 256
- IPv4 Routing: RIP, OSPF
- IPv4 Routes: up to 4,096
- OSPF Areas: up to 4

- OSPF Adjacencies: up to 64
- ECMP Paths: up to 4
- VRRP Instances: up to 256
- IPv6 Routing: Static
- Avaya VENA Distributed Top-of-Rack, technology:
 - Stack-mode of up to 8 units leveraging 5.12Tbps of virtual backplane capacity to support up to 256 10GbE ports, or up to 192 10GbE/16 40GbE ports
 - Fabric-mode of up to 500 units leveraging 280Tbps of virtual backplane capacity to support up to 16,000 10GbE ports, or up to 12,000 10GbE/1,000 40GbE ports

²10GBASE-T interfaces also support connection at 100/1000Mbps (Full Duplex); for both 7024XT and 7008XT.



- Avaya VENA Switch Cluster technology:
 - Standalone or Stacked mode
 - Triangle or Square configuration
 - up to 64 MLT Links
 - 128 SLT Links
 - VRRP Backup Master
 - SLPP, SLPP Guard
- Avaya VENA Fabric Connect, technology:
 - Standalone or Stacked mode
 - L2 Virtual Service Networks, L2 VSNs with Multicast¹, IP Shortcut Routing¹
 - up to 500 Nodes
 - 24 IS-IS Adjacencies
 - up to 500 Customer VLANs
 - 11,000 Service Identifiers
 - 4,096 Switched UNIs
- Software-Defined SAN/Fibre Channel over Ethernet
- Port Mirroring
- Remote Switch Port Analyzer (RSPAN)
- Unicast Storm Control

Pluggable Interfaces

- 40GBASE-SR4/4x10GBASE-SR QSFP+ up to 150m over MMF (MPO connector)
- 40GBASE-LR4 QSFP+ up to 400m over SMF (Duplex LC)
- 40GBASE-QSFP+ Direct Attach Cable, 1m
- 40GBASE-QSFP+ Direct Attach Cable, 3m

- 40GBASE-QSFP+ Direct Attach Cable, 5m
- 10GBASE-SR up to 300m reach over MMF (Duplex LC)
- 10GBASE-LRM up to 220m over FDDI-grade MMF (Duplex LC)
- 10GBASE-LR/LW up to 10km reach over SMF (Duplex LC)
- 10GBASE-ER/EW up to 40km reach over SMF (Duplex LC)
- 10GBASE-ZR/ZW up to 70km reach on SMF (Duplex LC)
- 10GBASE-ER CDWM up to 40km on SMF (Duplex LC)
- 10GBASE-ZR CDWM up to 70km on SMF (Duplex LC)
- 1000BASE-T up to 100m over CAT5E or better UTP Cable (RJ-45)
- 1000BASE-SX up to 550m reach on MMF (Duplex LC)
- 1000BASE-SX up to 550m reach on MMF (Duplex MTRJ)
- 1000-BASE-LX up to 550m reach on MMF, and up to 10 km on SMF (Duplex LC)
- 1000BASE-XD CDWM up to 40 km reach on SMF (Duplex LC)
- 1000BASE-ZX CDWM up to 70 km reach on SMF (Duplex LC)
- 1000BASE-EX up to 120 km reach on SMF (Duplex LC)
- 1000BASE-BX up to 10 and 40 km reach variants on SMF (LC)

IEEE and IETF Standards Compliance

- IEEE 802.1D Spanning Tree Protocol
- IEEE 802.1Q VLAN Tagging

- IEEE 802.1p Prioritizing
- IEEE 802.1s Multiple Spanning Tree
- IEEE 802.1w Rapid Reconfiguration of Spanning Tree
- IEEE 802.1AB Link Layer Discovery Protocol
- IEEE 802.1AX/802.3ad Link Aggregation Control Protocol
- IEEE 802.1ag Connectivity Fault Management, ITU-T Y.1731 Performance Management¹
- IEEE 802.1aq Shortest Path Bridging
- IEEE 802.1Qau Congestion Notification¹
- IEEE 802.1Qaz Enhanced Transmission Selection¹
- IEEE 802.1Qbb Priority-based Flow Control¹
- IEEE 802.3z Gigabit Ethernet over Fiber
 - 1 Gigabit, implemented as 1000BASE-SX, 1000BASE-ZX
- IEEE 802.3u Fast Ethernet
- IEEE 802.3x Flow Control
- IEEE 802.3z Gigabit Ethernet
- IEEE 802.3ab Gigabit Ethernet over Copper
- IEEE 802.3ae 10 Gigabit Ethernet over Fiber
 - 10 Gigabit, implemented as 10GBASE-SFP+: 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, 10GBASE-SW, 10GBASE-LW, 10GBASE-EW
- IEEE 802.3an 10 Gigabit Ethernet of Copper (10GBASE-T)
- 802.3ba 40 Gigabit and 100 Gigabit Ethernet over Copper and Fiber
 - 40 Gigabit, implemented as 40BASE-QSFP+: 40GBASE-SR4, 40GBASE-LR4
- RFC 768 UDP
- RFC 783 TFTP
- RFC 791/950 IP
- RFC 792 ICMP
- RFC 793 TCP
- RFC 826 ARP
- RFC 854 Telnet
- RFC 894 IP over Ethernet
- RFC 951 BootP
- RFC 1058 RIPv1
- RFC 1112 IGMPv1
- RFC 1157 SNMP
- RFC 1213 MIB-II
- RFC 1215 SNMP Traps Definition
- RFC 1271/2819 RMON

- RFC 1305 NTPv3
- RFC 1350 TFTP
- RFC 1361/1769 SNTP
- RFC 1493 Bridge MIB
- RFC 1573/2863 Interface MIB
- RFC 1583/2328 OSPFv2
- RFC 1643/2665 Ethernet MIB
- RFC 1757 RMON
- RFC 1850 OSPF v2 MIB
- RFC 1905/3416 SNMP
- RFC 1906/3417 SNMP Transport Mappings
- RFC 1907/3418 SNMP MIB
- RFC 1945 HTTP v1.0
- RFC 1981 Path MTU Discovery for IPv6
- RFC 2011 SNMPv2 MIB for IP
- RFC 2012 SNMPv2 MIB for TCP
- RFC 2013 SNMPv2 MIB for UDP
- RFC 2131 DHCP
- RFC 2138/2865/3576 RADIUS
- RFC 2236 IGMPv2
- RFC 2362 PIM-SM¹
- RFC 2460 Internet Protocol Version 6
- RFC 2474 DiffServ
- RFC 2475 DiffServ
- RFC 2548 Microsoft Vendor-specific RADIUS Attributes
- RFC 2674 Q-BRIDGE-MIB
- RFC 2737 Entity MIBv2
- RFC 2819 RMON MIB
- RFC 2865 Remote Access Dial-In User Service
- RFC 2866 RADIUS Accounting
- RFC 3046 DHCP Relay Agent Information Option
- RFC 3315 DHCPv6 (Relay)
- RFC 3376 IGMPv3
- RFC 3410 SNMPv3
- RFC 3411 SNMP Frameworks
- RFC 3412 SNMP Message Processing
- RFC 3413 SNMPv3 Applications
- RFC 3414 SNMPv3 USM
- RFC 3415 SNMPv3 VACM
- RFC 3584 Coexistence of SNMPv1/v2/v3
- RFC 3768 VRRP
- RFC 3917 IP Flow Information Export
- RFC 3954 NetFlow Services Export v9
- RFC 3993 Subscriber-ID Sub-option for DHCP¹
- RFC 4007 Scoped Address Architecture
- RFC 4022 TCP MIB
- RFC 4113 UDP MIB

- RFC 4213 Basic Transition Mechanisms for IPv6 (IPv6-in-IPv4)
- RFC 4293 IPv6 MIB
- RFC 4432 RSA Key Exchange for SSH
- RFC 4443 ICMPv6
- RFC 4861 Neighbor Discovery for IPv6
- RFC 4862 IPv6 Stateless Address Auto-Configuration
- RFC 5424 Syslog Protocol

Weights and Dimensions

- Height: 4.37cm, 1RU
- Width: 43.82cm
- Depth: 60.0cm
- Weight: 9kg for the base unit, and 12.2kg with Power Supplies, Fan Trays, and MDA installed

Power Specifications

- Input Voltage: 100-240 VAC
- Input Current
 - 7024XLS
 - 1.5-2.0 A @ 100-120 VAC
 - 0.75-1.0 A @ 200-240 VAC
 - 7024XT
 - 2.6-3.2 A @ 100-120 VAC
 - 1.3-1.6 A @ 200-240 VAC
- Power Consumption
 - 7024XLS
 - 180 W without MDA
 - 200 W with MDA
 - 400 W maximum
 - 7024XT
 - 240 W without MDA
 - 315 W with MDA
- Thermal Rating
 - 7024XLS – 615-1,366 BTU/h
 - 7024XT – 820-1,076 BTU/h

Environmental Specifications

- Operating temperature: 0 – 50°C
- Storage temperature: -40 to 85°C
- Operating humidity: 5 – 95% maximum relative humidity, non-condensing
- Storage humidity: 10 to 90% maximum relative humidity, non-condensing
- Operating altitude: 0 to 3,692 maximum
- Storage altitude: 0 to 12,192 maximum
- Acoustic Noise: less than 45 – 55dB at 35°C

RoHS Compliance

- Avaya Virtual Services Platform 7000 products, switches and field-replaceable components, are RoHS-compliant

Safety Agency Approvals

- Global basis for certification: EN 60950 current edition with CB national member deviations
- Mexico: complies with NOM

Electromagnetic Emissions & Immunity

- Global basis for certification: CISPR 22 Class A & CISPR 24, IEC 60950 with CB member national deviations
- US: complies with FCC CFR47 Part 15
- Canada: complies with ICES Class A
- Europe: complies with EN 55022 Class A; EN 55024; EN 300386 V1.3.3 Class A
- European Union & EFTA: complies with EN 55022; EN 55024; EN 61000-3-2; EN 61000-3-3
- Japan/Nippon: complies with VCCI
- Taiwan: complies with BSMI CNS 13428 & 14336, Class A
- Korea: complies with MIC Class A

Redundant Power

- 2 field-replaceable hot-swappable AC or DC internal Power Supplies

MTBF Values

- Avaya Virtual Services Platform 7000 base unit: 241,000 hours

Warranty

- Lifetime Next Business Day hardware replacement
- Lifetime Basic Technical Support
- 90-Day Advanced Technical Support
- Optional Software Release Service also available: GW5300ASG / GW6300ASG

Country of Origin

- China (PRC)

About Avaya

Avaya is a leading, global provider of customer and team engagement solutions and services available in a variety of flexible on-premise and cloud deployment options. Avaya's fabric-based networking solutions help simplify and accelerate the deployment of business critical applications and services. For more information, please visit www.avaya.com.

© 2015 Avaya Inc. All Rights Reserved.

Avaya and the Avaya logo are trademarks of Avaya Inc. and are registered in the United States and other countries. All other trademarks identified by ®, TM, or SM are registered marks, trademarks, and service marks, respectively, of Avaya Inc.
04/15 • LB5289-08