Ethernet Switching

Engineering

Super Large Campus
Technical Configuration Guide

Avaya Data Solutions

Document Date: June 2011

Document Number: NN48500-609

Document Version: 1.4

Notices
While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer
Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer
Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation(s) provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty
Avaya provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product, while under warranty, is available to Avaya customers and other parties through the Avaya Support Web site: http://www.avaya.com/support
Please note that if you acquired the product from an authorized reseller, the warranty is provided to you by said reseller and not by Avaya.

Licenses

Copyright

Third Party Components
Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: http://support.avaya.com/Copyright.

Trademarks
The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the documentation(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party. Avaya is a registered trademark of Avaya Inc. All non-Avaya trademarks are the property of their respective owners.

Downloading documents
For the most current versions of documentation, see the Avaya Support. Web site: http://www.avaya.com/support

Contact Avaya Support
Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: http:// www.avaya.com/support.

Abstract

This Technical Solution Guide defines the recommended designs for a Super Large Converged Campus infrastructure. The document provides an overview of the best design practices to implement a network capable of supporting converged applications and services.

The audience for this Technical Solution Guide is intended to be Avaya Sales teams, Partner Sales teams and end-user customers. All of these groups can benefit from understanding the common design practices and recommended components for a converged campus network design.

For any comments, edits, corrections, or general feedback, please contact Dan DeBacker (ddebacke@Avaya.com).

Table of Contents

Figures

Tables

Conventions

This section describes the text, image, and command conventions used in this document.

Symbols

👍 Tip – Highlights a configuration or technical tip.

ⓘ Note – Highlights important information to the reader.

△ Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

Text

Bold text indicates emphasis.

Italic text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

```
ERS5520-48T# show running-config
```

Output examples from Avaya devices are displayed in a Lucida Console font:

```
ERS5520-48T# show sys-info

Operation Mode:        Switch
MAC Address:           00-12-83-93-B0-00
PoE Module FW:         6370.4
Reset Count:           83
Last Reset Type:       Management Factory Reset
Power Status:          Primary Power
Autotopology:          Enabled
Pluggable Port 45:     None
Pluggable Port 46:     None
Pluggable Port 47:     None
Pluggable Port 48:     None
Base Unit Selection:   Non-base unit using rear-panel switch
sysDescr:              Ethernet Routing Switch 5520-48T-PWR
                       HW:02       FW:6.0.0.10  SW:v6.2.0.009
Mfg Date:12042004    HW Dev:H/W rev.02
```

1. Converged Campus Design Solutions

The Converged Campus architecture is built using the fundamental strategic values of the Avaya Data Solutions organization. By adhering to these core values, Avaya provides a solid infrastructure on which the enterprise can build upon. With this solid infrastructure, the enterprise can solve their business challenges by enabling services easily and without worry. Avaya offers a unique value proposition in its ability to provide this infrastructure while still offering best-in-class total cost of ownership.

Figure 1.1: Avaya Data Solutions Strategic Values

The Converged Campus solutions have been broken down into Small, Medium, Large, and Super Large to address specific requirements of the Enterprise. A major objective of these Technical Solution Guides is to provide a blueprint and starting point for the customer network design. By providing solutions that have been architected, validated, and documented, the building block for the network is now in place and ready for the specific customization required by each individual network. This customization comes in the form of specific VLANs required, protocols being used, number of edges to connect, and application requirements for the infrastructure.

This solution guide provides optimal network designs and general best practices when implementing and administering the network. The end result is a network that can sustain both normal data traffic as well as any converged applications deployed in the enterprise.

ⓘ Note – All design recommendations and best practices within this guide should be reviewed against the available features on the Ethernet switching platforms being deployed and should also be reviewed against the release notes for the versions of software being used. As feature enhancements are introduced and bugs are fixed, it is imperative to understand the capabilities and limitations of the switches and software being implemented. This ensures that the design being deployed utilizes the features and functions of the switches to their maximum effectiveness.

1.1 Avaya Converged Enterprise

The architecture shown in Figure 1.2 includes all areas of the Avaya Converged Enterprise solution. This guide focuses specifically on the Converged Campus architecture for edge switching and core switching. There are many permutations of possible designs when deploying infrastructure from Avaya, but this guide highlights the major design concepts that need to be addressed.

The ultimate goal of these designs is to provide a highly reliable infrastructure with sub-second seamless failover preventing any interruption of traffic on the network. The value in this is two-fold. First, in the event of a failure, no loss of connectivity or traffic will be experienced by the end user. Secondly, and probably just as important, is the ability to provide near hitless software upgrades for the core of the network.

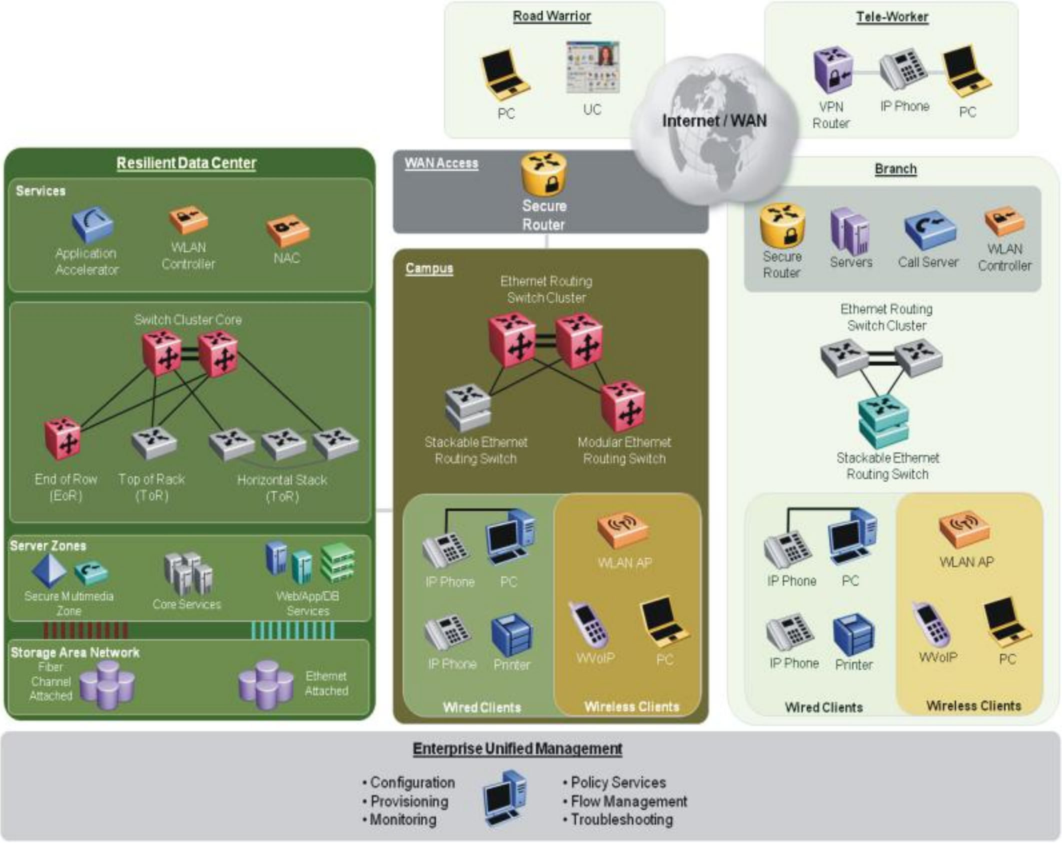

Figure 1.2: Converged Enterprise Architecture

1.2 Chassis versus Stackable

Several factors come into play when choosing the edge switching solution. Consider the following criteria when selecting the edge product while keeping in mind that the stacking technology continues to evolve and is getting closer and closer to simulating a modular chassis solution in many respects.

Switch reliability is a key concern. In the past, modular switches were thought to be more reliable with redundant power supplies, redundant fan trays, and redundant switch fabrics and CPUs. However, the evolution of the stackable switch has reduced the disparity between the two platforms by employing a resilient stacking architecture, supporting internal or external redundant power supplies, and providing features such as auto unit replacement and new unit quick configuration. Both solutions can provide an equally highly reliable edge solution today.

Scalability of the edge switch includes the ability to add ports easily, increase bandwidth out of the closet, and add protocol and features within the closet. A chassis solution typically adds ports by adding new input/output (I/O) modules in the chassis, while stackable switches add ports by adding switches to the existing stack. Both solutions limit the total number of ports supported in a single stack/chassis. The stackable switches provide more flexibility when adding bandwidth out of the closet. A stack can be broken up into two or more stacks, thus increasing bandwidth out of the closet very easily.

As stackable switches are added to the closet, each one must be powered individually, which uses several outlets in the closet. In contrast, only two to four outlets are usually required for a chassis.

The same protocols and features are for the most part available on both platforms; however, scalability of those protocols is normally greater in a chassis solution. It is easier to redeploy stackable switches as a stack or standalone unit, whereas the modular chassis requires additional hardware to support the I/O modules.

Serviceability and manageability differences between the two solutions are minimal. With both solutions, you can add ports easily, perform software upgrades, retain multiple configurations, and manage the stack or chassis as a single entity.

Rack space can also be a consideration when selecting the edge-switching platform. Typically, a stackable solution takes up less total rack space than a chassis solution in both height and depth. However, stackable switches require rear access for power connections and stacking connections, whereas a chassis solution requires only front access.

The final consideration between the two solutions is price. Usually, a chassis solution is slightly more expensive than a stackable solution due to the additional Switch Fabric/CPU (SF/CPU), chassis, and power supplies needed. In summary, both solutions offer great reliability and scalability. Customers must decide which provides the optimal solution for their organization.

1.3 Layer 2 versus Layer 3 at the Edge

The process of choosing between Layer 2 and Layer 3 at the edge can take many different twists. When considering the differences between the two, it is imperative to keep in mind the end goal of 99.999 percent network availability. There are several ways to design a Converged Campus network. The goal is to design a network that provides high reliability, fast convergence, and yields the lowest possible total cost of ownership (TCO). The TCO is derived by adding the initial cost of equipment/installation (CAPEX – Capital Expenditures) and the ongoing administration and support of the network (OPEX – Operating Expenditures). Over the long run, the OPEX is often higher than the CAPEX, so the goal is to help reduce OPEX by making the network easy to administer and troubleshoot.

The two major areas to consider when deciding between Layer 2 and Layer 3 at the edge are (1) IP routing and (2) intelligence, which can be thought of as operating at Layers 3 to 7. Intelligence can further be defined as the ability to provide traffic management (QoS and content-aware switching) and security, which includes end user authentication and policy enforcement. The goal is to centralize the routing and distribute the intelligence to provide a high-performing and secure network along with easy and simplified management. However, one must also consider the number of users being aggregated. The ability to distribute ARP tables across the network may prove a more efficient design. There are no absolute numbers to tell you whether your network should centralize all routing or distribute the routing, but guidelines are provided in the design recommendations below.

A Layer 2 edge solution, when combined with strong distributed intelligence features, is easier to implement, administer, and troubleshoot. In addition, sub-second failover and no penalties on performance make Layer 2 the clear superior choice in the Converged Campus design.

Avaya, however, recognizes that a Layer 2 solution is not always possible or may not fit every network design. The Avaya edge switch portfolio includes products that support a Layer 3 edge into a Layer 3 core/distribution. There are no performance penalties for implementing Layer 3 at the edge. The switches provide outstanding performance whether implemented as Layer 2 or Layer 3. The main difference is seen in the complexity of laying out the Layer 3 design and the ongoing administration and troubleshooting of such a network.

In summary, Avaya provides the flexibility for both approaches. Some customers choose a Layer 3 edge design solution for various reasons – no VLAN propagation, same configuration replicated, smaller broadcast domains, security/access control lists (ACL), for example – and they have the necessary routing expertise to support such a network. Other customers prefer a centralized routing and filtering/ACL approach, which may reduce the overall complexity of the network administration by not distributing Layer 3 throughout the network.

2. Super Large Campus Design

The Super Large Campus design is intended to support highly scaled networks in terms of the number of supported devices. In the example provided here, the network contained 10000 network devices as a lower limit, with the platform capable of supporting much greater numbers. Please take note that these numbers are network attached devices such as PCs, IP phones, printers, access points, etc. and ***not*** users. Attempting to base a network design on users is becoming increasingly difficult as more devices are being converged onto the infrastructure; therefore, recommendations are based on network-attached devices.

The ***Super Large Campus Solution*** includes the following key components along with design and best practice recommendations for:

- Virtual Services Platform 9000 (VSP 9000) at the Core
- Ethernet Routing Switch 8800 (ERS 8800) at the Distribution Layer, if present
- Ethernet Routing Switch 2500, 4500, 5000 Series, and 8300 at the Edge
- Identity Engine Ignition Servers for Network Access Control
- Media Gateway for VoIP services

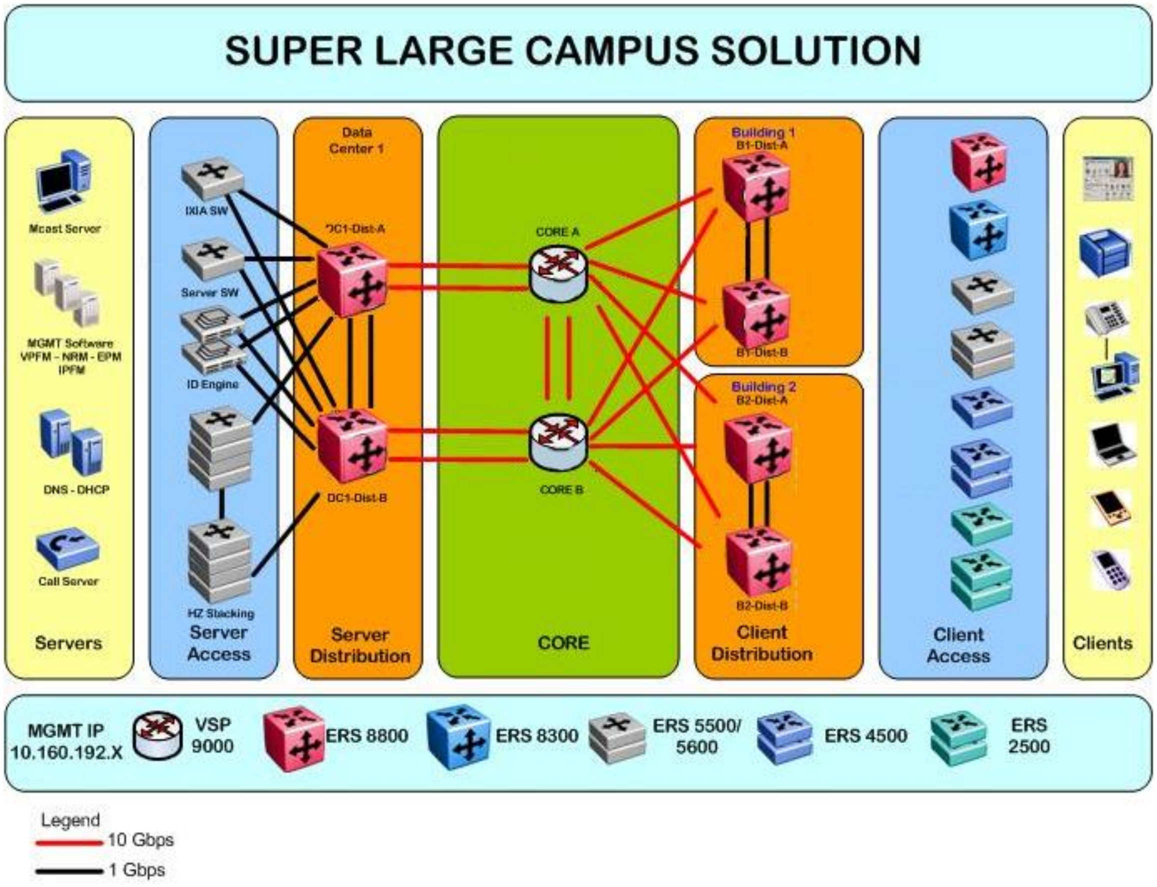

Figure 2.1: Super Large Campus Ethernet Infrastructure

Avaya took the following business requirements into consideration when selecting the products used in this specific solution:

- Effective and Efficient Edge Switching
- Scalability
- Cost-effective Without Compromising Performance
- High Availability
  - Edge with Resilient Stacking
  - Switch Clustering Core
- Simple to Build and Run
- Energy Efficiency

The following figure show an example of a two-tiered architecture where the Edge connects directly to the Core using 10GE connections.

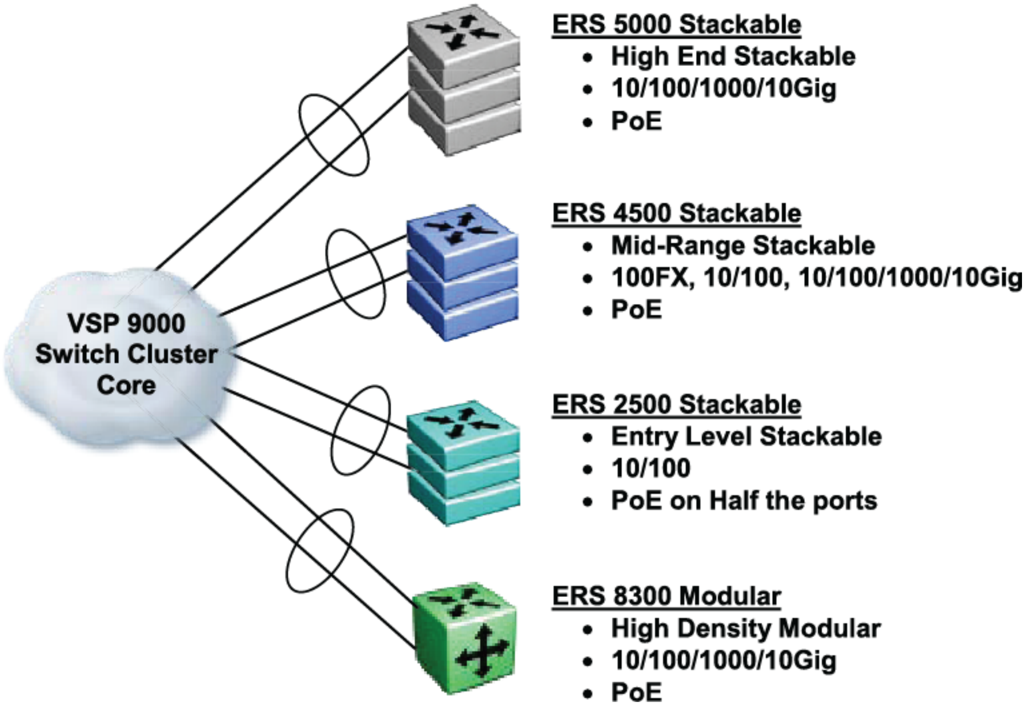

Figure 2.2: Two-Tiered Topology Example

The following figure shows an example of a three-tiered architecture where there is a Distribution Layer between the Edge and the Core. The connections from the Edge to the ERS 8800 in the Distribution Layer are GE; the connections from the Distribution Layer to the Core are 10GE.

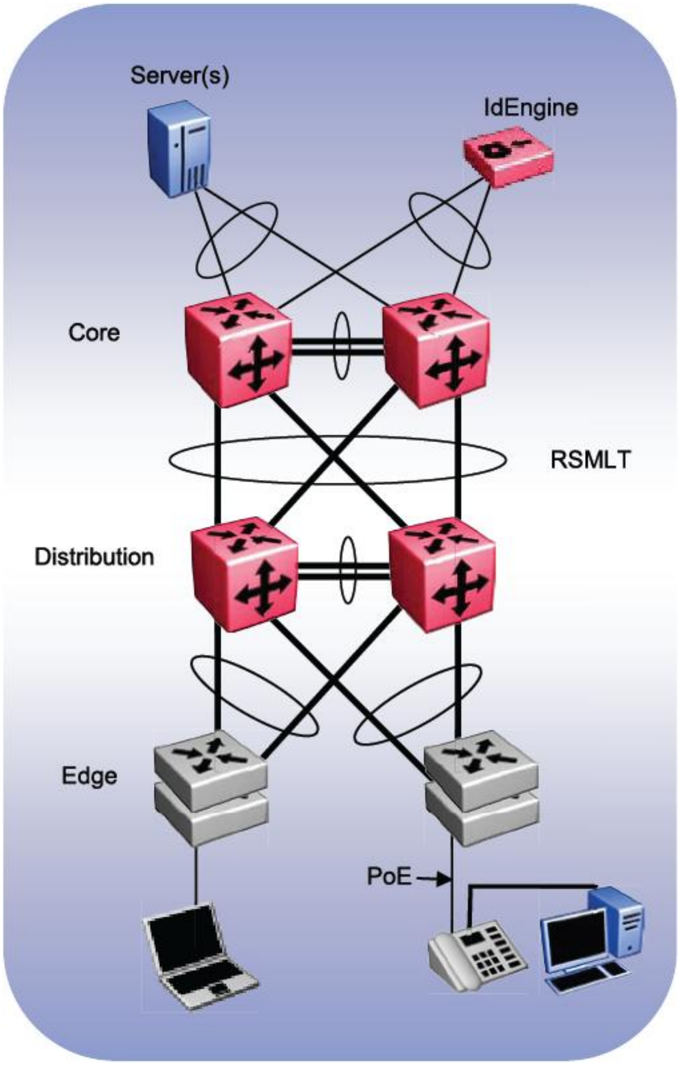

Figure 2.3: Three-Tiered Topology Example

2.1 Core Switching

The VSP 9000 platform serves as the Core switching platform for the Super Large Campus solution. The features and functionality highlighted here represent the basic requirements for the Super Large Campus. Please refer to the product documentation for a detailed explanation of these and all the other features of the VSP 9000 platform.

- Core Switching Hardware
- Advanced Software License
- Switch Clustering
- VLACP
- SLPP
- Filter Untagged Frames
- Spanning Tree
- VLANs
- DHCP Relay
- Quality of Service
- Layer 3
- VRRP with Backup Master
- Server Connectivity

2.1.1 Core Switching Hardware

The VSP 9000 is a chassis-based solution with several I/O module options to fit the need of the Super Large Campus Core and Data Centers. The VSP 9000 initially supports a switching architecture of 8.4 Tbps that will scale to 27 Tbps in a single chassis and over 100 Tbps in a quad Switch Cluster. The VSP 9000 supports high density configurations with up to 480 Gigabit ports (copper or fiber) or up to 240 10 Gigabit ports in a single chassis offering the flexibility needed for most core switching solutions. The VSP 9000 also provides significant investment protection with support for 40 Gigabit and 100 Gigabit in the future along with its flexible packet processor that enables new software capabilities without a need to upgrade hardware.

The VSP 9000 9024XL module has 24 ports and each continuous physical group of 4 ports supports a combined bandwidth of more than 11Gbps. The 9024XL module also has a 3.5:1 oversubscribed line rate over 24 ports; 6 ports can provide full line rate if you do not use the remaining ports. You can use these 6 ports (marked with a black square around the port number) to achieve full line rate for the attached interfaces. Use only a single port from each grouping to ensure no oversubscription at 10Gbps, and in addition one port out of each group can be used with a 1Gbps SFP and still be line rate .

The VSP 9000 also supports advanced serviceability features such as flash memory on cards for Flight Recorder capture for enhanced after crash debugging, Key Health Indicators for instantenous health monitoring, run-time diagnostics, detailed packet counters in chips to isolate packet failures, and checksums to detect packet data corruption.

The following figure shows the front and rear view of the VSP 9000 chassis. The front slots support two control processor modules and ten I/O modules; the rear slots support two auxiliary modules (for future use) and six switch fabric modules.

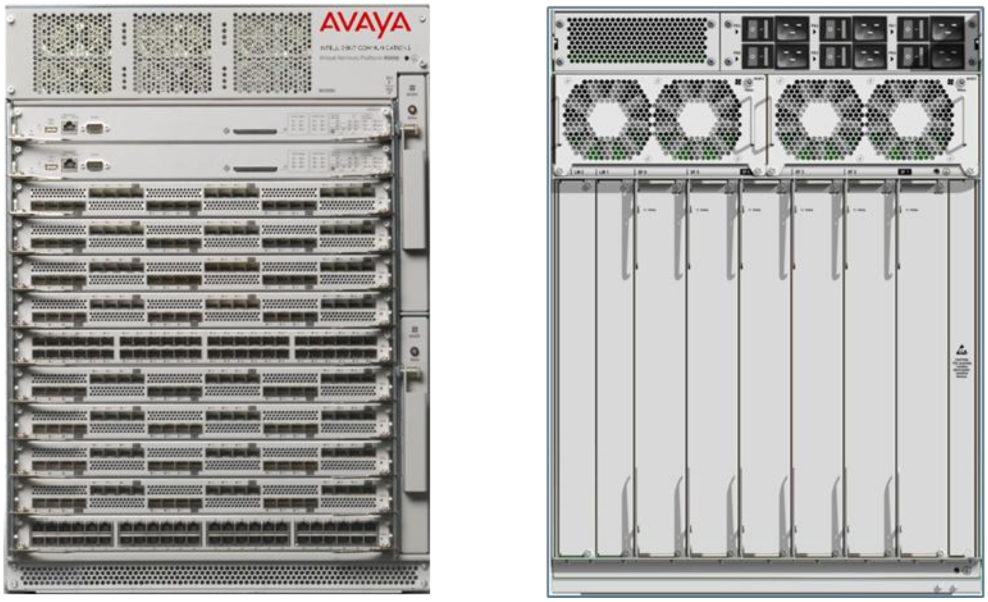

Figure 2.4: VSP 9000 Chassis

The following table lists the VSP 9000 modules.

Module	Ports	Type
9024XL	24	24 ports of SFP / SFP+ supporting 1GbE and 10GbE transceivers
9048GB	48	48 ports of SFP supporting 100M and 1GbE transceivers
9048GT	48	48 ports supporting 10/100/1000
9090SF	N/A	Switch Fabric
9080CP	N/A	Control Processor (CPU module)

Table 2.1: VSP 9000 Modules

The VSP 9000 I/O modules support a variety of pluggable for both Gigabit and 10 Gigabit.

- Gigabit – SX, LX, XD, ZX, BX, CWDM, and copper
- 10 Gigabit – SR, LR, ER, ZR, LRM

2.1.2 Dual-CPU High Availability

The VSP 9000 supports High Availability (HA) with dual CP modules and Rapid Failure Detection and Recovery technology. The dual 9080CPs provide support the following features:

- 1+1 control plane redundancy
- Rapid Failure Detection and Recovery (RFDR) of Data Path -- <20ms
- Hardware-assisted control plane overload protection
- Hardware-assisted traffic profile-based DOS detection to protect control plane
- Redundant GE Chassis Area Network (CAN)
- Hitless patching
- LEDs that duplicate the LEDs of the modules in the back of the chassis

You can configure the CPUs to operate in either a HA mode or non-HA mode. If you want to switch from one mode to the other, you only have to reboot the standby CP. The master CP does not require rebooting.

You can operate the dual CPUs in either HA or non-HA mode. The default mode is HA enabled.

- In **HA mode**, also called hot standby, the platform synchronizes the two CPUs. The CPUs use the same configuration and forwarding tables. The master CPU automatically updates the forwarding tables of the secondary CPU in real time. If the master CPU fails, the secondary CPU takes over the master responsibility very quickly, thereby minimizing traffic interruption for the failure condition.
    - When there is a switchover in hot standby mode, the standby CP becomes the master CP, the master CP resets, and the I/O and SF modules continue to run.
- In **non-HA mode**, also called warm standby, the platform does not synchronize the two CPUs. If the master CPU fails, the secondary CPU must start before it can take on the master responsibility, and then must also relearn the forwarding table information. This operation causes an interruption to traffic.

  When there is a switchover in warm standby mode, the VSP 9000 resets all the I/O and SF modules, reloads the configuration file, and relearns all the tables.

The VSP 9000 provides *full* HA application support to the following:

- ARP
- DHCP Relay
- EAP
- ECMP
- IGMP L2 Snooping
- LACP/VLACP
- MAC, Unknown MAC, Global MAC Filter, and MAC Learning
- MLT/ RSMLT
- NLB
- OSPF
- RIP
- Routing Policies
- RSTP/MSTP
- SLPP
- SMLT
- SONMP
- Static route
- UDP broadcast forwarding
- VLAN
- VRF-Lite
- VRRP

The VSP 9000 provides *partial* HA application support to the following:

- BGP
- IGMP (with PIM enabled)
- IPv6
- PIM-SM
- PIM-SSM

For more details, please refer to the ***VSP 9000 Administration Guide (NN46250-600)***.

2.1.3 High Availability Core

The core switching layer is the most critical component of the Super Large Campus design. To ensure that the Super Large Campus has the most redundant and resilient system possible, Avaya recommends a resilient switch cluster of VSP 9000s each with dual 9080 Control Processor (CP) modules, six 9090 Switch Fabric (SF) modules, and N+1 redundant power supplies and fans. All components are hot swappable and fully redundant with no single point of failure.

The 9080CP runs Avaya's carrier-grade Linux Operating System to provide a reliable and secure architecture. The 9080CP supports all high level protocols, and distributes the results (routing updates) to the rest of the system, manages and configures the interface and SF modules, and maintains and monitors the health of the chassis.

The 9090SF modules provide the cross-bar switching solution in the midplane chassis. Each 9090SF module connects to ten different interface modules and two CP modules simultaneously. Each chassis supports up to six SF modules and can be configured to support N+1 switch fabric redundancy.

2.1.4 Hitless patching

A patch is a small piece of software designed to fix and correct a problem in software code, its supporting data, or update a feature. The VSP 9000 delivers bug fixes and feature enhancements in a faster and simpler fashion than a full software upgrade. It can also be a vehicle to deliver debug code to gather valuable information within the network.

There are two types of patching supported in this release:

- Hitless Patch

  The main feature of hitless patching is that the target software is patched without any disruption to running system processes. As a general rule, hitless patches will be the primary mode of patch creation. Other patching options will be considered only if all options to create a hitless patch have been exhausted.

- Reset Patch

  Reset patches are provided as an alternative form of patching when a hitless patch cannot be created as a fix for a problem. After reset patches are applied to the target software, the operator must reset the system for the patches to take effect. There will be a service outage and disruption of traffic when the system is reset.

Before applying a patch to a running system, make sure that you read the README file that came with the patch. For additional information, see ***Avaya Virtual Services Platform 9000 Upgrades and patches — Software Release 3.0 (NN46250-400)***.

2.1.5 Flight Recorder

The Flight Recorder is a framework in place on the VSP 9000 that stores both history and current state information for various kernel, system, and application data with minimal overhead to execution. This data can later be accessed on demand when debugging system issues to give engineers the best possible troubleshooting information. Functionally, the Flight Recorder consists of two elements: Persistent Memory and Always-on Trace.

The Persistent Memory feature stores information in volatile memory that persists across processor resets. This feature provides information on crashes, errors, and outages that are not the result of a power failure. Persistent Memory data not saved to non-volatile storage before a power failure will be lost. Persistent Memory snapshots are taken when:

- a critical process stops functioning
- a card resets
- the hardware watchdog activates
- the user initiates a snapshot in the ACLI

The Always-on Trace feature creates an ongoing, circular log of every trace call recently executed regardless of the trace level enabled by the user. This functionality provides storage for central processor trace records, storage for input/output trace records, and storage for switch fabric trace records. Since the Always-On Trace performs circular logging, reading the log from top to bottom will not represent a chronological sequence of events. Pay attention to timestamp information to discern the chronology of events.

Flight Recorder functionality is provided only through the ACLI. For more information on Flight Recorder, see ***Avaya Virtual Services Platform 9000 — Troubleshooting (NN46250-700)***.

2.1.6 Key Health Indicators

The Key Health Indicators (KHI) feature of the Avaya Virtual Services Platform 9000 provides a subset of health information that allows for quick assessment of the overall operational state of the device. The KHI feature is not intended to provide a comprehensive debugging solution. Instead, KHI identifies key information that could lead a technician towards discovery of a specific failure. After KHI information has been assessed, further debugging will be required to determine the specific reason for the fault. Avaya recommends capturing KHI information during normal operations to provide a baseline for detecting fault situations.

KHI information is provided only through the ACLI. For more information on KHI, see ***Avaya Virtual Services Platform 9000 — Fault Management (NN46250-703)***.

2.1.7 Advanced and Premier Software License

The VSP 9000 offers three different licenses: the Base License (required with every chassis), the Advanced License, and the Premier License. The Premier License encompasses all features within the Premier category as well as the Advanced License category. There is no upgrade license available to go from Advanced to Premier.

Enabling features on a Virtual Services Platform 9000 requires the generation and installation of a license file that contains the authorized MAC addresses of the switches on which you install the license file.

You must purchase one Base software license for each chassis. The optional Advanced and Premier feature licenses provide access to additional features contained within those licensing levels. Purchase these licenses separately in the form of either an Advanced License Kit or Premier License Kit. The Premier License Kit contains all Advanced License Kit features. When you purchase either an Advanced

License Kit or a Premier License Kit, the license covers all current and future features. If you currently have an Advanced License Kit, no discounted price exists to move to a Premier License Kit; you must purchase a complete Premier License Kit. If you purchase a Premier License Kit, you have licenses for all features for the life of the product. For more information, contact your Avaya sales representative.

Base (BA)	Advanced (AD)	Premier (PR)
▪ Core Layer 2 Switching, ACLs, policers, shapers, 802.1D/w/s, 802.1p/Q ▪ MLT/LACP/SMLT/RSMLT/SLT ▪ Core Layer 3 routing & switching (IGMP, PIM, RIPv1/v2, OSPF) ▪ VRRP, SLPP, IPv6 Mgt, DoS protection	▪ Base features plus: ▪ BGP4 (16 peers / 64k IP Routes) ▪ Packet Capture function (PCAP)	▪ Advanced features plus: ▪ 500k IP Routes data plane ▪ 1.5M BGP control plane routes ▪ 256 BGP peers ▪ VRF

Table 2.1.7 – Virtual Services Platform 9000 3.0 Licensing Levels

Tip – Avaya recommends that you purchase the Premier License if you anticipate growth in your network. If you purchase the Advanced License, and later require features available only if you have the Premier License, you must also purchase the Premier License. If you purchase the Premier License initially, you have access to all features enabled by the Advanced License and the Premier License (you do not need to purchase the Advanced License separately).

You must purchase the Base software license for each chassis. You can install an Advanced or Premier License on each chassis after you install the Base software license, but the Advanced and Premier Licenses are optional.

2.1.7.1 Premier Trial License

The Virtual Services Platform 9000 provides a trial period of 60 days during which you have access to all features. In the trial period you can configure all features without restriction, including system console and log messages. This trial period starts when the chassis is initially configured with the software image.

For additional information on software licensing please reference ***Avaya Virtual Services Platform 9000 – Administration (NN46250-600)***.

2.1.8 Two Tier versus Three Tier Architecture

When designing a Converged Campus solution, there are two major topologies that can be implemented. The two-tier architecture, in which all edge switches terminate into the core of the network, and the three tier architecture, in which the edge switches terminate into a distribution layer network. The distribution layer network then terminates into the core. Three-tier architecture is usually required when the existing cable plant cannot support a two-tier deployment because of fiber distances or physical layout of the fiber.

From a switching/routing perspective, there are two options to be considered: either Layer 2 at the edge with Layer 3 in the core/distribution, or Layer 3 at the edge with Layer 3 in the core/distribution. Avaya provides Ethernet switching platforms that can provide either design alternative. There is no right answer for all possible designs; however, the Avaya design philosophy is always to keep the architecture simple without compromising resiliency and scalability. This translates into easier management and an overall lower total cost of ownership (TCO) by centralizing routing in the core and distributing intelligence across the network.

Avaya recommends deploying two-tier architecture whenever possible. This simplifies the network, reduces the amount of equipment required, and does not compromise scalability and resiliency. The two-tier architecture supports either Layer 2 or Layer 3 at the edge. Avaya recommends keeping Layer 2 VLANs at the edge and routing between VLANs at the core.

If three-tier architecture is deployed, Avaya recommends using Layer 3 between the distribution and core layers, utilizing RSMLT when possible for these connections. The same rules apply to the connections between the access and distribution layers (for less than 3000 users, use Layer 2; for more than 3000 users, use Layer 3).

With any of these options, it is critical to deploy an end-to-end QoS strategy to ensure that mission-critical applications are able to provide the required quality of experience for the users. A detailed discussion of QoS is covered in a later section of this document.

In a two-tiered architecture, the Edge connects directly to the Core using 10GE connections. In a three-tiered architecture where there is a Distribution Layer between the Edge and the Core, the connections from the Edge to the Distribution Layer are GE and the connections from the Distribution Layer to the Core are 10GE.

2.1.9 Two Tier Design – Core to Edge

With the basic two-tier design, the Edge switches connect directly into the Core. In the Super Large Campus design, the core is a Resilient Switch Cluster consisting of a minimum of two VSP 9000s with sufficient port density to accommodate dual homing of all edge switches.

An Inter-Switch Trunk (IST) ties the pair of VSP 9000s together to form the Resilient Switch Cluster. The IST is a critical component of the Switch Cluster and therefore must be highly resilient. The architecture of Switch Clustering and the traffic flow through the cluster is such that there is not a high volume of traffic across the IST, so resiliency of the connection is more important than total bandwidth.

The architecture is very flexible and can accommodate most design scenarios. The standard recommendation is to have Layer 2 at the Access and Layer 3 at the Core. This architecture does not preclude the ability to extend Layer 3 to the Access if that is desired.

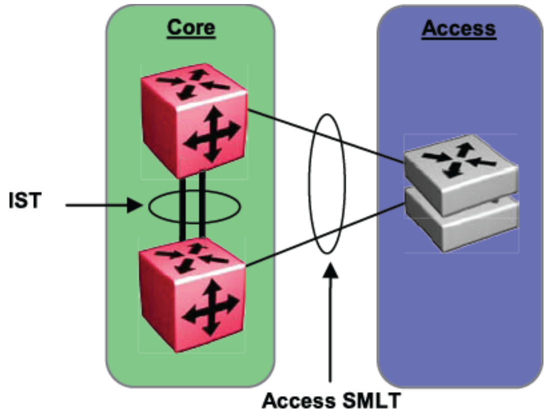

Figure 2.5: Switch Clustering – Two Tier Architecture

2.1.10 Three Tier Design – Core to Distribution to Edge

With a three-tier design, a VSP 9000 or ERS 8800 Distribution layer is inserted between the Edge and the VSP 9000 Core. The need for a distribution layer can be attributed to existing physical infrastructure such as fiber plant layout or the requirement to connect multiple buildings on a single campus together – in that case, the building cores would be the distribution layer and connect back to a centralized core.

In situations where three tiers are necessary, there are options on the deployment of Switch Clustering for resiliency between the layers. The edge to distribution considerations are the same as described in the above section. Between the distribution and core layers, there are different options available based on the architecture deployed:

- Layer 2 between Distribution and Core – SMLT

  In the attempt to centralize routing functionality and distribute the intelligence throughout the network, it is easy to keep a simple Layer 2 architecture between these two layers of the network. In this design, the distribution to core connectivity mimics that of the edge to the core described in the above section. The main difference lies in the ability to fully mesh the distribution to the core. A fully meshed solution provides the highest level of resiliency possible with still maintaining sub-second failover and recovery. A square or full mesh is mandatory to maintain full resiliency and bandwidth between distribution and core.

- Layer 3 between Distribution and Core – Routed SMLT

  If routing is desired between the distribution and core layers, deploy routed SMLT (RSMLT) to maintain sub-second failover and recovery while running a standard IGP routing protocol such as RIP or OSPF. RSMLT builds on the SMLT technology by providing an active-active router concept to SMLT networks with routing enabled on the core VLANs. In the case of a routing switch failure, RSMLT takes care of packet forwarding at Layer 2 while the routing protocol converges at the Layer 3 level. This allows the non-stop forwarding of traffic in the event of any failure with no disruption to the user. Another huge advantage of RSMLT is the ability to extend Layer 2 subnets – something that is not possible if strictly using Layer 3 routing between the core and distribution.

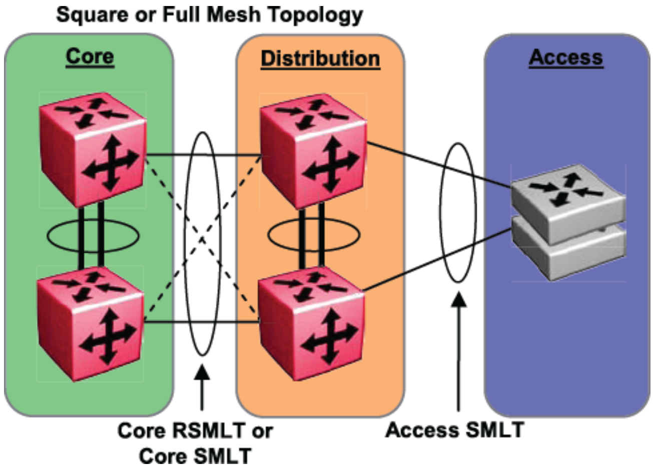

Figure 2.6: Switch Clustering – Three Tier Architecture

2.1.11 Switch Clustering using Split MultiLink Trunking (SMLT)

Switch Clustering using Split MultiLink Trunking (SMLT) provides industry-leading technology for the resiliency of the Converged Campus design. Providing redundant links that forwards traffic with no Spanning Tree allows the ultimate design in a converged environment. Sub-second failover and the simplicity of a network without Spanning Tree reduce TCO and ensure converged applications will function flawlessly. A vital feature of Switch Clustering is its ability to work with any end device (3rd party switch, servers, etc.) that supports a form of link aggregation.

Switch Clustering also provides the ability to perform virtual hitless upgrades of the core switches (cluster). With all connections to the cluster dually attached, a single core switch can be taken out of service without interrupting end user traffic. This switch then can be upgraded and brought back into service. By performing the same function on the other switch, after the upgraded switch is back online, the entire cluster can be upgraded without taking a service outage and with minimal interruption to traffic flows on the network.

Switch Clustering supports 512 groups and shares the number of MLT IDs across the cluster. This includes MultiLink Trunking and Split MultiLink Trunking (SMLT). It also supports 480 trunks with up to 16 links per group.

2.1.12 Switch Clustering Terminology

There are different design options to consider with the deployment of Switch Clustering:

➢ Single Link Trunking (SLT)

> Note – SLT is not available on the VSP 9000 the system support 512 MLT ID's thus use normal MLT/SMLT instead of SLT. However, if you have a three-tiered topology with a Distribution Layer in the Core, the ERS 8800 does support SLT.

SLT is a port-based option allowing large-scale deployments of SLT from a single Switch Cluster. Every port (saving at least two for the IST) can be used for SLT groups terminating into the cluster, with each SLT group consisting of a maximum of two uplinks (one per core Ethernet Routing Switch). For most typical deployments, the ability to have two connections per edge switch/stack is more than sufficient bandwidth, and allows a single cluster to handle many environments. The flexibility of the Avaya edge switch solutions allows for uplinks ranging from 10 Mbps to 10 Gbps (uplinks within the same SLT group must be of the same media type and link speed).

➢ Split MultiLink Trunking (SMLT)

The MLT-based SMLT option allows for increased scaling of the number of links within a single SMLT group. The number of links supported in an SMLT group is the same number of MLT links supported on the Ethernet Routing Switch platform being used for the Switch Cluster. The SMLT links can be spread across the Switch Cluster – usually in an even dispersion, but this is not an absolute requirement. One MLT group must be used to create the IST between the two switches used to form the Switch Cluster.

Both SLT and SMLT can be configured on the same Switch Cluster.

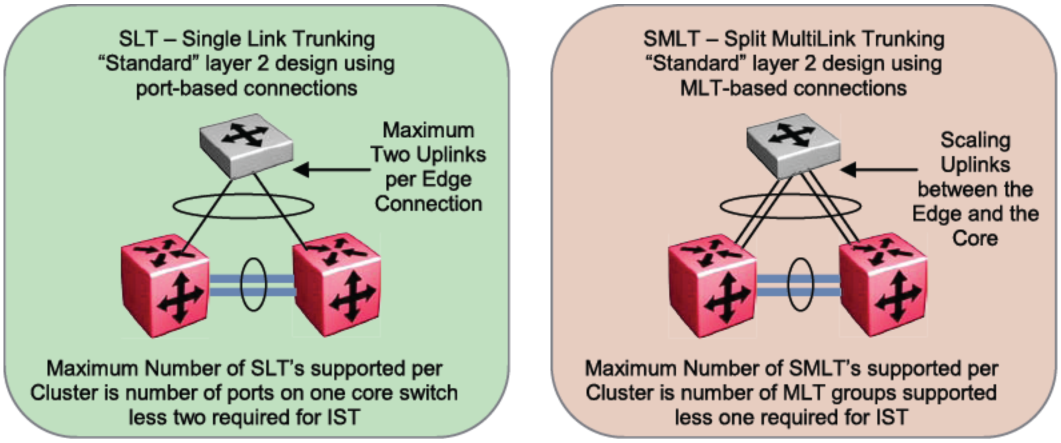
SLT – Single Link Trunking
"Standard" layer 2 design using
port-based connections
Maximum
Two Uplinks
per Edge
Connection
Maximum Number of SLT's supported per
Cluster is number of ports on one core switch
less two required for IST
SMLT – Split MultiLink Trunking
"Standard" layer 2 design using
MLT-based connections
Scaling
Uplinks
between the
Edge and the
Core
Maximum Number of SMLT's supported per
Cluster is number of MLT groups supported
less one required for IST

Figure 2.7: SLT and SMLT Terminology

Switch Model	Links per MLT Group	MLT Groups per Switch or Stack	MLT-based SMLT Groups			Port-based SLT Groups		
			Copper	**Fiber (1GbE)**	**Fiber (10GbE)**	**Copper**	**Fiber (1GbE)**	**Fiber (10GbE)**
VSP 9000	16	512*	511*	511*	511*	N/A	N/A	N/A
ERS 8800	8	128	127	127	127	382	238	22
ERS 8300	4	31	30	30	30	382	398	67
ERS 5000	8	32	31	31	31	398	190	62

* The number of MLT groups that the VSP 9000 supports is limited to the number of ports provided by the I/O modules installed in the chassis. In Release 3.0, the maximum number of 1GbE ports is 480, and the maximum number of 10GbE ports is 240.

Table 2.2: MLT/SMLT/SLT Scaling Capabilities

Note – An Advanced Software License is required for the ERS 8300 and ERS 5000 for Switch Clustering (SMLT/SLT).

2.1.13 Switch Clustering Topologies

There are three supported topologies with Switch Clustering. Choose the best topology based on the overall design of the network.

- Triangle – Single Switch Cluster at the core with the edge directly connected via SMLT
- Square – Two pairs of Switch Clusters interconnected by SMLT. Squares can be scaled with additional pairs of Switch Clusters
- Full Mesh – Expanding on the Square topology, the full mesh adds additional connections between the pairs so that each switch has at least one connection to every other switch in the square. Full Mesh topologies can be scaled with additional pairs of Switch Clusters.

The following sections highlight the supported topologies that you can use with the VSP 9000 as the Switch Cluster core.

2.1.13.1 Triangle Switch Cluster

The triangle Switch Cluster is comprised of a single VSP 9000 switch for each IST peer. This configuration terminates Edge closet connections using MLT-based SMLT.

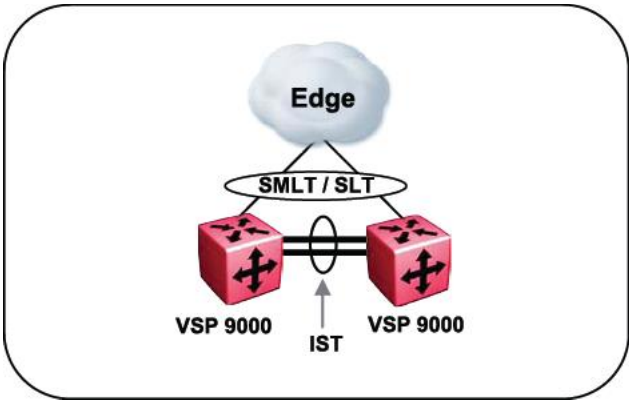

Figure 2.8: Triangle Switch Cluster

2.1.13.2 Square / Full Mesh Switch Cluster

The square or full mesh Switch Cluster extends the scalability of the VSP 9000 by connecting Switch Cluster Cores using SMLT or RSMLT. This configuration terminates Edge closet connections using MLT-based SMLT. All rules from triangle configurations still apply. In this topology, the system performance can exceed 100 Tbps.

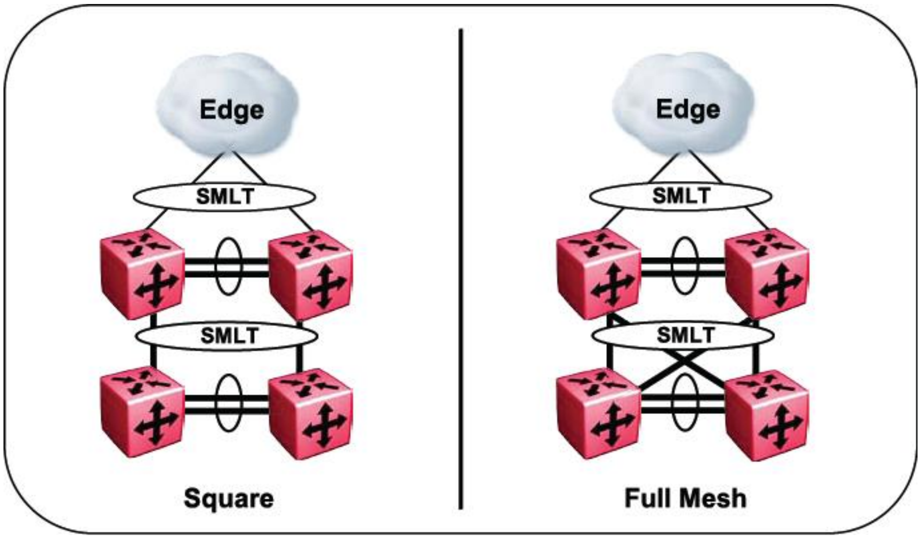

Figure 2.9: Square / Full Mesh Switch Cluster Topologies

The square or full mesh Switch Cluster provides the basic infrastructure to overlay various logical designs. Whether running all Layer 2, all Layer 3, or a combination of Layer 2 on one Switch Cluster and Layer 3 on the other Switch Cluster, this design accommodates those needs.

2.1.14 Switch Clustering Reference Architecture

In order to easily identify different aspects of the Switch Cluster design, the following reference architecture will be used throughout the discussion on best practice recommendations. The following diagram depicts a six-switch core for completeness, showing the triangle, square and full mesh topologies. Please note that this is not a requirement for implementing Switch Clustering.

- Access SLT/SMLTs are connections from the core out the edge closets and are normally in a standard triangle configuration.
- Core SMLT/RSMLT connections exist between Switch Clusters and can be formed using either the square or full mesh topologies. SMLT connections are used for the core so that bandwidth can easily be increased by adding another connection to the MLT group that forms the SMLT.

The major difference between the Access and the Core will be in the Loop Prevention mechanisms recommended for each. The Core is obviously more critical to the overall network and is also a much more controlled environment; therefore, the best practice implementation will differ between the Access and the Core. A more detailed discussion on these techniques will follow in the coming sections.

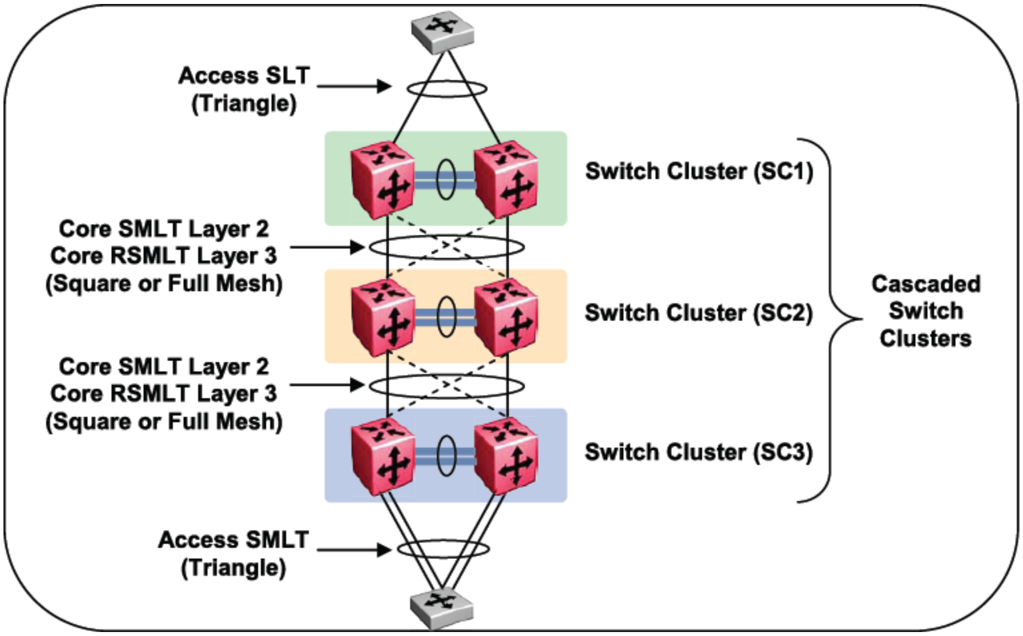

Figure 2.10: Switch Clustering Reference Architecture

2.1.15 Switch Cluster Core Configuration Guidelines

When configuring Switch Clustering on the VSP 9000 review the following:

- IST is 2-port MLT (minimum) and up to 16-port MLT (maximum)
  - The number of links required for the IST is based on the amount of bandwidth required during a failure scenario. The amount of traffic that traverses the IST during normal operations is minimized with the Switch Clustering architecture and all connected devices being dual-homed.
  - Ports within the MLT must be same speed.
  - Mixed media MLTs (copper and fiber) are supported for the IST.
  - Avaya strongly recommends using Distributed MLT (DMLT) connections between modules in the chassis for added resiliency.
  - Ports assigned to an MLT (IST) are indexed by a number starting at zero (0). The lowest port position (module 1 port 1) for an MLT link is assigned an index of zero. The next MLT link in the second lowest position gets an index of one (1) and so on. This index is used by the MLT algorithm to assign a flow over a particular MLT link. Therefore, Avaya recommends that you mate the lowest port position of one MLT link in one switch with the lowest port position of the peer switch. Follow this rule for all successive MLT links. This will help to ensure that the MLT algorithm always resolves a flow over the same link between the two switches.
  - Avaya recommends that the IST links terminate on non-blocking ports in each chassis whenever possible. This ensures that all critical management traffic is received properly by each switch regardless of the utilization of the IST links. At the same time, having the IST comprised of oversubscribed ports is supported and is sometimes the only option available. With multiple ports comprising the IST and the fact that multiple failures would likely be required to overrun the IST, these designs will work in most cases. However, it is important to understand potential limitations of such designs.
  - Do not use the IST IP addresses as next hop addresses for any static routes.
  - Both IPv4 and IPv6 SMLT is supported, however the IST only supports IPv4 messaging.
- If you have an ERS 8800 Distribution Layer and **only two uplinks** are required per edge closet, use port-based Single Link Trunking (SLT). The number of configurable SLTs equals the number of ports in the switch less the number of ports which are used for the IST.
  - Begin SLT ID’s at 129
- If you have an ERS 8800 Distribution Layer and **more than two uplinks** are required per edge closet, use MLT-based SMLT. The number of configurable SMLTs is 128. One MLT-based SMLT will be used for the IST.
  - Use MLT 1 for the IST.
  - SMLT ID’s 2 to 128 should correspond with MLT ID – although this is not required it is highly recommended to simplify the design and troubleshooting in the future.

- If you do not have an ERS 8800 Distribution Layer and more than two uplinks are required from the edge closet to the VSP 9000 Core, use MLT-based SMLT. The number of configurable SMLTs is 512. One MLT-based SMLT will be used for the IST.
  - Use MLT 1 for the IST.
  - SMLT ID's 2 to 512 should correspond with MLT ID – although this is not required it is highly recommended to simplify the design and troubleshooting in the future.
  - You can configure both SLT and SMLT connections on the same ERS 8800 Switch Cluster simultaneously. On VSP 9000 Switch Clusters, you can configure SMLT only.
  - It is possible to overlap the ID numbers when using MLT-based SMLT and port-based SLT, Avaya recommends avoiding this and following the recommendations in Table 2.3.

Switch Model	Software Version	MLT-based SMLT ID's	Port-based SLT ID's
VSP 9000	3.0 and higher	1-512	Not Supported
ERS 8800	4.1 and higher	1-128	129-512
ERS 8300	3.0 and higher	1-31	32-512
ERS 5000	5.0 and higher	1-32	33-512

Table 2.3: SMLT ID Recommended Values

- Create a separate IST VLAN and use a private address space for the IST VLAN IP addresses with a small subnet mask (i.e. 30-bit mask). This VLAN is only required for IP communications between IST peers.
- Verify that all VLANs participating in SLT/SMLT are configured on both IST peer switches and are tagged on both ends of the IST.
- All SLT/SMLT uplinks shall be 802.1Q tagged, as this will easily facilitate adding additional VLANs to the edge without impacting traffic. This will also ensure that any switches added to the network in default configuration will not cause loops, as untagged frames will be discarded at the Switch Cluster core – see details that follow in the section on Filter Untagged Frames.
- STP will automatically be disabled on all ports participating in SLT/SMLT, this includes both the IST and SLT/SMLT ports. However STP must be manually disabled on the edge switch ports that connect to the SLT/SMLT ports.
- If multicast routing (PIM-SM) is enabled on the Switch Cluster (VSP 9000 and ERS 8800), enable PIM-SM on the IST VLAN to insure fast recovery of multicast traffic.
- When configuring a Core SMLT square or full mesh (SMLT between two pairs of Switch Clusters), use the same SMLT ID on both sides of the square/mesh for operational simplification.

2.1.16 VLANs

VLANs provide an easy mechanism for traffic separation, a way to minimize the size of broadcast domains, and can help isolate different protocols from each other. In most cases, the VLAN is considered equal to a broadcast domain; for example, a specific IPv4 or IPv6 subnet is assigned to a single VLAN. From an administrative point of view, VLAN to subnet mapping makes each very easy to identify quickly. The number of VLANs and the type of VLANs deployed can vary greatly from design to design. Consider these points when creating the VLAN strategy:

- The need for isolation of different protocols on the network
- VLAN types to be used – port, protocol, MAC, subnet
- Traffic separation for voice and data
- Size of the broadcast domain/number of users
- VLANs by geographic area – per closet, per floor, per building
- Network services required per VLAN – DHCP, UDP forwarding, etc.

All Voice and Data VLANs will be configured on both core VSP 9000 switches. Any other services that will be attached directly to the Switch Cluster core will also need VLANs as well as the IST.

- By default, VLAN 1 is created and all ports are members – this VLAN cannot be deleted. Do not use this VLAN for any production traffic, but only as a repository for unused ports.
- Increase the FDB timer on all Switch Cluster VLANs from the default of 300 seconds to 21601 seconds (1 second greater than the ARP timer) for the VSP 9000. This reduces the amount of re-ARPs that need to occur.
- Any VLAN with an IP Address can be used to manage the switch; however, to simplify management of the switch, create a management VLAN. As an alternative you may also use the out-of-band management port on the CP Module for management.

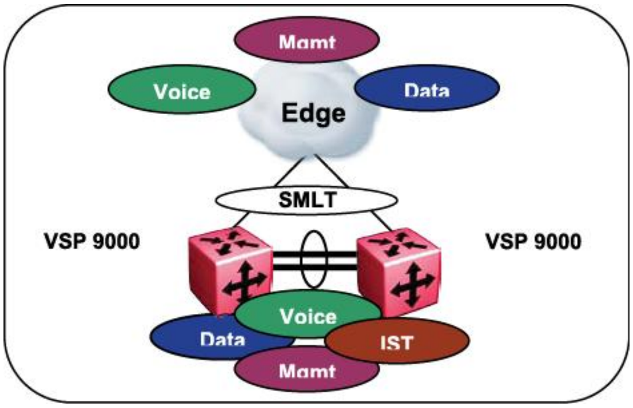

Figure 2.11: VSP 9000 VLANs

The following table shows the order of precedence for identifying packets entering the switch.

Highest Priority	→				Lowest Priority
Hardware Platform	**802.1Q Tagged Packet**	**Subnet Based VLAN**	**Protocol Based VLAN**	**MAC Based VLAN**	**Port Based VLAN**
VSP 9000	✓	✓	✓	✓	✓
ERS 8800	✓	✓	✓	✓	✓
ERS 8300	✓	N/A	✓	N/A	✓
ERS 5000	✓	N/A	✓	N/A	✓

Table 2.4: VLAN Support

2.1.17 Discard Untagged Frames

The recommendation to always enable 802.1Q VLAN tagging on uplink ports is critical to using the discard untagged frames feature. This provides protection against a factory defaulted or incorrectly configured device being connected to the VSP 9000 or ERS 8800. The core will automatically drop all packets that are not 802.1Q tagged, adding another level of protection against potential loops.

Tip – The Discard Untagged Frames feature should be enabled on the IST/SLT/SMLT ports.

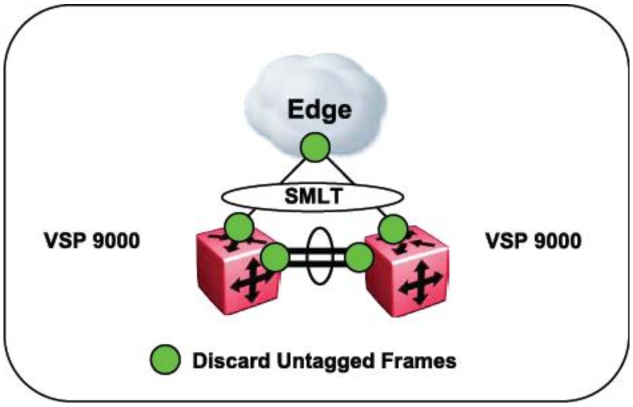

Figure 2.12: VSP 9000 Discard Untagged Frames

2.1.18 Spanning Tree

Avaya recommends using the Spanning Tree protocol on all end station connections in order to safeguard the network from hubs or other devices that could be inserted into the network at the end station. For more information about the edge connections, see section 2.2.9 Spanning Tree Protocol.

In the VSP 9000, the default spanning tree protocol is MSTP. However, Spanning Tree is automatically disabled on all IST / SMLT ports on the VSP 9000 Switch Cluster Core.

Note – When using SMLT to connect the edge to the distribution/core, always disable Spanning Tree on the uplink ports/MLT of the edge switch.

Avaya recommends using Split MultiLink Trunking (SMLT) to interconnect closets to the core of the network, thus eliminating the need for the Spanning Tree protocol on uplinks. In addition to automatically disabling Spanning Tree on all IST / SMLT ports on the VSP 9000 Switch Cluster Core, Spanning Tree is also automatically disabled on all IST / SLT / SMLT ports on the ERS 8800 in order for Switch Clustering to function properly. Avaya recommends leaving all the other ports as spanning tree enabled to ensure protection from a hub or other network device causing a loop in the core.

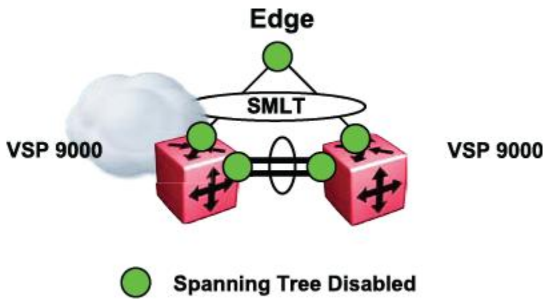

Figure 2.13: VSP 9000 Spanning Tree

2.1.19 Control Plane Rate Limit (CP-Limit)

Control plane rate limit (CP-Limit) controls the amount of control traffic that can be sent to the CP from a physical port (i.e. OSPF hello, RIP update, etc.). Packets exceeding the configured packet rate value are dropped. In this way, CP-Limit protects the CP from being flooded by traffic from a single, unstable port. This differs from normal port rate limiting that limits non-control multicast traffic and non-control broadcast traffic on the physical port, which would not be sent to the CP (i.e. IP subnet broadcast, etc.). Configure the CP-Limit feature on a per-port basis within the chassis.

The CP-Limit default settings are:

- Default state is enabled on all ports
- When creating the IST, CP-Limit is disabled automatically on the IST ports

If the actual rate of packets-per-second sent from a port exceeds the defined rate, then the port is administratively shut down to protect the CP from continued bombardment. An SNMP trap and a log file entry are generated indicating the physical port that has been shut down as well as the packet rate causing the shut down.

To reactivate the port, you must first administratively disable the port and then re-enable the port. Alternatively, you can enable AutoRecoverPort. This feature activates an automatic recovery of the port from the action taken by CP Limit. The AutoRecoverPort feature is disabled by default.

Having CP-Limit disable IST ports in this way could impair network traffic flow, as this is a critical port for SMLT configurations. Avaya recommends that an IST MLT contain at least 2 physical ports, although this is not a requirement. Avaya also recommends that CP-Limit be disabled on all physical ports that are members of an IST MLT – this is the default configuration. Disabling CP-Limit on IST MLT ports forces another, less-critical port to be disabled if the defined CP-Limits are exceeded. In doing so, you preserve network stability should a protection condition (CP-Limit) arise. Although it is likely that one of the SMLT MLT ports (risers) would be disabled in such a condition, traffic would continue to flow uninterrupted through the remaining SMLT ports.

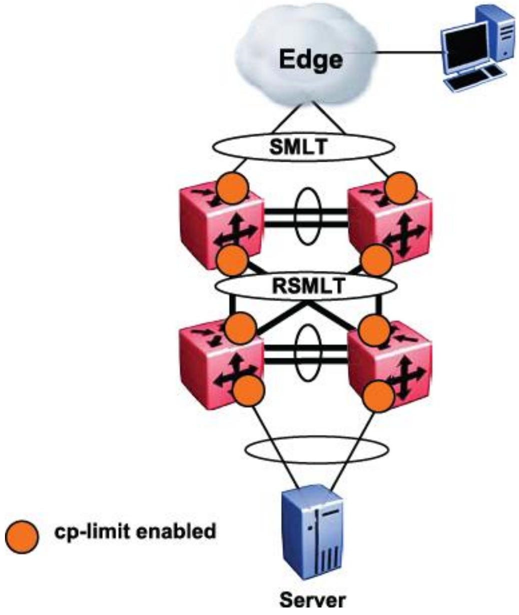

Figure 2.14: CP-Limit Recommendations

- Leave enabled on all SLT/SMLT ports in the distribution/core.
- Disable for all ports in the IST. Ports participating in the IST should never be shut down under any circumstances.
- In multi-tiered core environments, Avaya recommends that edge closet switches have CP-Limit values less than the values used on the core links. This way, if an offending device does transmit malicious traffic, the edge switches will get triggered because of lower values, thus preventing the important core links from shutting down. This will also aid in isolating problems.
- CP-Limit works only on ports that belong to VLANs that have Layer 3 enabled.

For edge and server connected ports, if the connected device is determined to produce traffic to the levels for which cp-limit is configured, the connected port will be disabled when it starts transmitting.

Recommended CP-Limit Values		
	Broadcast	Multicast
Aggressive		
Workstation	1000	1000
Server	2500	2500
Non-IST Interconnection	7500	7500
Moderate		
Workstation	2500	2500
Server	5000	5000
Non-IST Interconnection	9000	9000
Relaxed		
Workstation	4000	4000
Server	7000	7000
Non-IST Interconnection	10000	10000

Table 2.5: CP-Limit Recommended Values

△ Caution – Altering CP-Limit values from their defaults during normal network operation can cause the links to become disabled. Avaya strongly recommends that to obtain a baseline of the network traffic across the links, choose the right value, and apply.

2.1.20 VLACP

Avaya developed Virtual LACP (VLACP) to provide a true end-to-end failure detection mechanism between directly connected switches or connectivity across intermediary networks. This feature now adds a greater level of resiliency and flexibility to the Converged Campus design when used in conjunction with MLT, DMLT, and SMLT.

Please note that LACP and VLACP are two totally independent features and one does not require the other for implementation. LACP provides standards-based link aggregation capabilities and point-to-point failure detection, while VLACP provides end-to-end failure detection only.

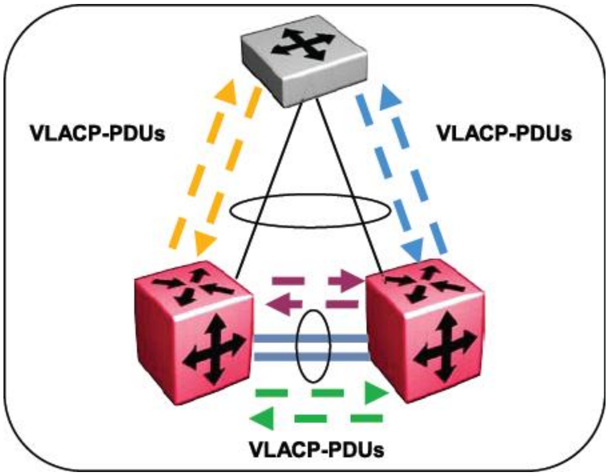

Figure 2.15: Virtual Link Aggregation Control Protocol (VLACP)

The following highlights the features associated with VLACP:

- Designed to operate end-to-end, regardless of whether the switches are directly connected or have an intermediary connection between them.
- VLACP is strictly a heart-beat end-to-end detection mechanism with no link aggregation capabilities.
- Global and port-based configuration parameters
- Very light load on CPU processing
- On each port that has VLACP enabled, VLACP PDUs are sent periodically. If VLACP PDUs are not received on a particular link, that link is administratively disabled after a configurable timeout period.
- Can run independently as a port-to-port protocol or on top of MLT, DMLT, and SMLT.

VLACP is a critical feature when deploying resilient networks with Switch Clustering. VLACP can detect end-to-end failures as described above and can also disable links to a switch that might still have link connectivity but cannot process traffic due to unexpected switch lockups. VLACP can also prevent loops in the network if uplinks are plugged into the wrong ports; VLACP is enabled only on uplinks, therefore, when the uplink is plugged into the wrong port without VLACP enabled, the link at the other end will be administratively disabled and no traffic will use that link.

When VLACP is used on directly connected switches, Avaya recommends using a reserved multicast MAC address for the VLACP PDU. This reserved MAC address will not be flooded or forwarded by a switch that receives it. In the case where a factory-defaulted switch is connected, the VLACP packets would not be flooded back down to the core (giving a false positive that the links should be brought up), thus ensuring integrity of the network.

Enabling VLACP on the IST is critical for Switch Cluster designs. This will protect the integrity of the core in the case where one of the switches was inadvertently reset to factory default. Without VLACP on the IST, data packets are sent to the defaulted switch, causing loops and various other types of forwarding issues. By enabling VLACP on the IST, these packets would not be sent across, as the ports would be administratively down. Also, by having VLACP on all the SMLT/SLT ports, the edge switches would not send or receive packets from the defaulted switch, thus somewhat alleviating a much larger problem.

In the Super Large Campus design, VLACP will be used between the VSP 9000 Switch Cluster and all uplinks to the ERS 2500/4500/5000/8300 switches at the edge as well as on the IST.

- Globally configure VLACP to use the reserved multicast MAC of 01-80-C2-00-00-0F
- For the IST links, use the long timeout
    - Slow periodic timer of 10000 msecs * timeout scale of 3

Make sure these values match on both ends of the links

- For the SLT/SMLT links, use the short timeout
- Fast periodic timer of 500 msecs * timeout scale of 5
    - Make sure these values match on both ends of the links

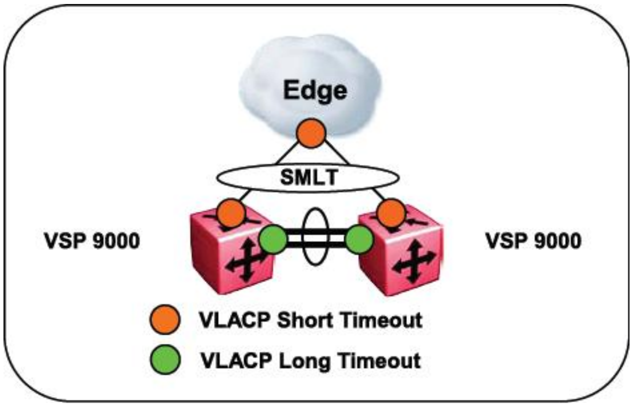

Figure 2.16: VSP 9000 VLACP

2.1.21 Simple Loop Prevention Protocol (SLPP)

Simple Loop Prevention Protocol (SLPP) provides active protection against Layer 2 network loops on a per-VLAN basis. SLPP uses a lightweight hello packet mechanism to detect network loops. SLPP packets are sent using Layer 2 multicast and a switch will only look at its own SLPP packets or at its peer SLPP packets. It will ignore SLPP packets from other parts of the network. Sending hello packets on a per VLAN basis allows SLPP to detect VLAN based network loops for un-tagged as well as tagged IEEE 802.1Q VLAN link configurations. Once a loop is detected, the port is shutdown. The SLPP functionality is configured using the following criteria:

- SLPP TX Process – the network administrator decides on which VLANs a switch should send SLPP hello packets. The packets are then replicated out all ports which are members of the SLPP-enabled VLAN. Avaya recommends enabling SLPP on all VLANs.
- SLPP RX Process – the network administrator decides on which ports the switch should act when receiving an SLPP packet that is sent by the same switch or by its SMLT peer. You should enable this process only on Access SMLT/SLT ports and never on IST ports or Core SMLT/SLT ports in the case of a square/full mesh core design.
- SLPP Action – the action operationally disables the ports receiving the SLPP packet. The administrator can also tune the network failure behavior by choosing how many SLPP packets need to be received before a switch starts taking an action. These values need to be staggered to avoid edge switch isolation – see the recommendations at the end of this section.

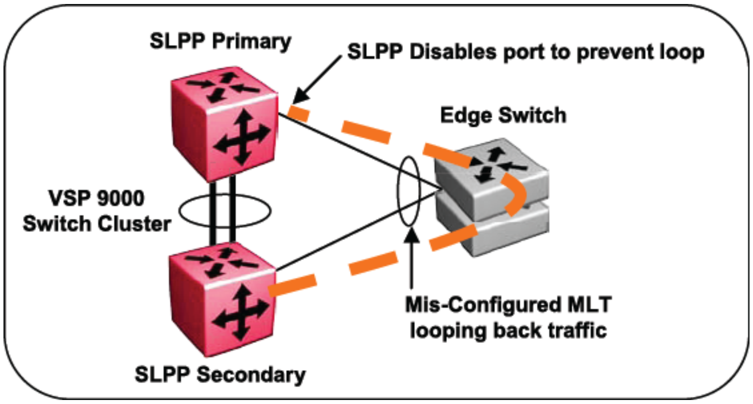

Figure 2.17: Simple Loop Prevention Protocol (SLPP)

Loops can be introduced into the network in many ways. One way is through the loss of an MLT configuration caused by user error or malfunctioning equipment. This scenario may not always introduce a broadcast storm but, because all MAC addresses are learned through the looping ports, it does significantly impact Layer 2 MAC learning. Spanning Tree would not in all cases be able to detect such a configuration issue, whereas SLPP reacts and disables the malfunctioning links, limiting network impact to a minimum.

The desire is to prevent a loop from causing network problems while also attempting to not totally isolate the edge where the loop was detected. Total edge closet isolation is the last resort in order to protect the rest of the network from the loop. With this in mind, Avaya adopted the concept of an SLPP Primary switch and SLPP Secondary switch. These are strictly design terms and are not configuration parameters. The Rx thresholds are staggered between the primary and secondary switch, therefore the primary switch will disable an uplink immediately upon a loop occurring. If this resolves the loop issue, the edge closet still has connectivity back through the SLPP secondary switch. If the loop is not resolved, the SLPP secondary switch will disable the uplink and isolate the closet to protect the rest of the network from the loop.

As the number of VLANs running SLPP scale off of a specific uplink port, the Rx-threshold value may need to be increased to prevent complete isolation of the offending edge. Critical to note is that the primary goal of SLPP is to protect the core at all costs. In certain loop conditions, what may occur is the secondary switch also detecting the loop and its SLPP Rx-threshold is reached before the primary can stop the loop by taking its port down. Therefore, both switches eventually take their ports down and the edge becomes isolated. The larger the number of VLANs associated with the port, the more likely this could occur, especially for loop conditions that affect all VLANs.

To accommodate different design requirements, the following table details the different options for SLPP RxThreshold values for edge connections:

Recommended SLPP RxThreshold Values		
Threshold	**Primary**	**Secondary**
Aggressive	5	50
Moderate	2 times the number of VLANs Minimum SLPP value = 5 Maximum SLPP value = 100	10 times the number of VLANs Minimum SLPP value = 50 Maximum SLPP value = 500
Relaxed	3 times the number of VLANs Minimum SLPP value = 5 Maximum SLPP value = 100	15 times the number of VLANs Minimum SLPP value = 50 Maximum SLPP value = 500

Table 2.6: SLPP Recommended Values – Access Edge

In the Super Large Campus design, SLPP will be configured on the VSP 9000 Switch Cluster(s) as described below:

- Designate one VSP 9000 core switch as primary and the other as secondary. This is strictly from a design perspective. There are no actual configuration parameters to create the primary or secondary.
- Enable SLPP globally and on all VLANs on the Switch Cluster **except** for IST VLAN
- SLPP Rx should never be configured on IST ports
- Enable SLPP Rx (threshold of 5) on all non-SMLT ports in the VSP 9000 except the IST ports as added protection against possible loops in the core
- Do not enable auto recovery – once the port is disabled by SLPP, it will need to be re-enabled manually after the loop has been fixed
- SLPP Rx-threshold is NOT reset upon any activity, but is a cumulative count. This can cause a situation in which multiple different loop events, can lead to an event where both primary and secondary links have their threshold reached and both links bring their ports down, and edge isolation could occur. A **disable/enable of SLPP**, which does not impact the network, should be performed after any SLPP event to clear the counters.

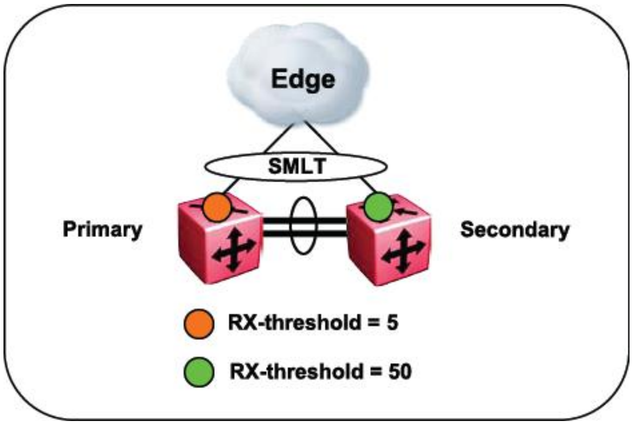

Figure 2.18: VSP 9000 SLPP in Triangle Topology

For Square and Full Mesh configurations that use a bridged core (Layer 2 VLANs extend from the edge through all switches in the core), Avaya recommends enabling SLPP on the primary switch as shown below. Enabling SLPP on half of the core will still prevent any possible loops and will not allow the possibility of the entire core being shut down by a loop at the edge of the network.

Recommended SLPP RxThreshold Values			
	🟠	🟢	🟣
Aggressive	5	50	300
Moderate	2 times the number of VLANs (5-100)	10 times the number of VLANs (50-500)	400
Relaxed	3 times the number of VLANs (5-100)	15 times the number of VLANs (50-500)	500

Table 2.7: SLPP Recommended Values – Bridged Core

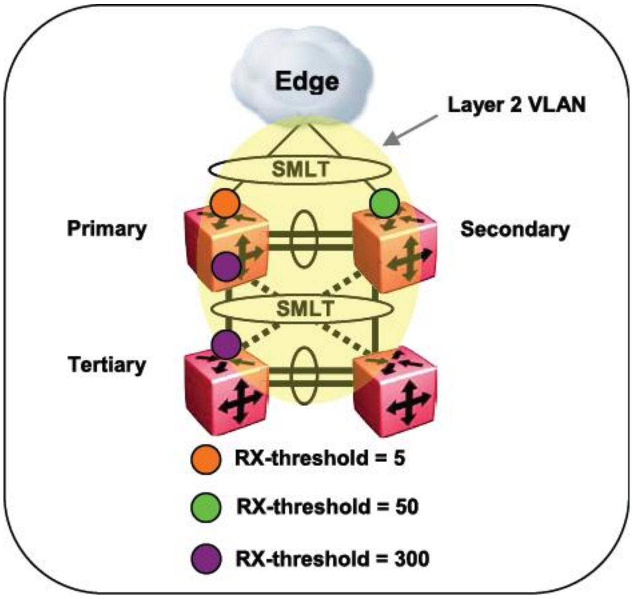

Figure 2.19: SLPP in Square/Full Mesh Bridged Core

For Square and Full Mesh configurations that use a routed core, Avaya recommends using or creating a separate core VLAN and enabling SLPP on that VLAN and the square or full mesh links between the Switch Clusters. Loops created in the core will be caught and loops at the edge will not affect core ports. If using RSMLT between the Switch Clusters, then enable SLPP on the RSMLT VLAN.

Since SLPP will only be enabled on one or two VLANs in the core, changing the RX-threshold values will not likely be necessary.

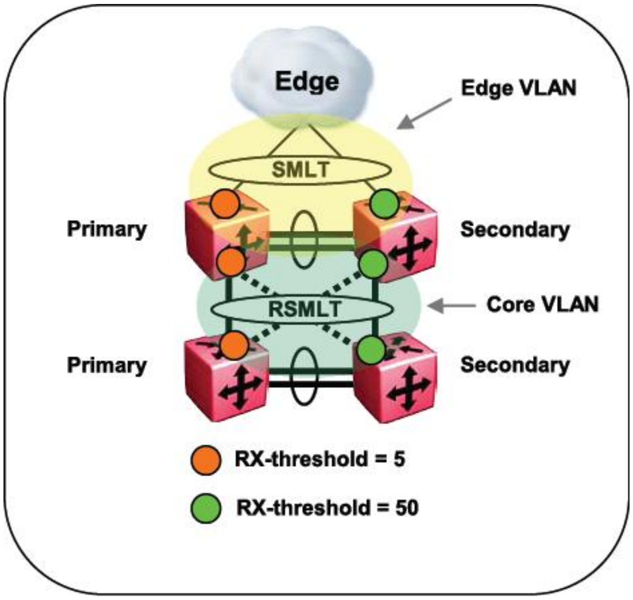

Figure 2.20: SLPP in Square/Full Mesh Routed Core

2.1.22 Quality of Service

Differentiated Services (DiffServ) is the industry and Avaya standard for the implementation of QoS. DiffServ provides QoS on a per hop behavior in the Ethernet switches by marking the header of individual packets with a DiffServ Code Point (DSCP). This DSCP then provides an indication to the Ethernet switch as to the priority of each packet and into which queue the packet should be placed. Note that the network infrastructure must support QoS end to end. Without a full end-to-end deployment, QoS cannot provide the necessary actions to ensure priority through every hop of the network. By default, the edge switches re-mark all QoS bits to zero and do not honor any markings.

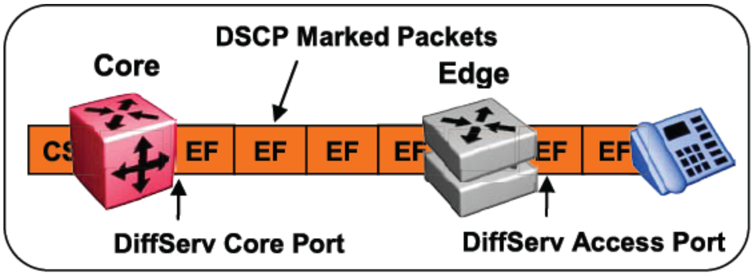

Figure 2.21: Quality of Service

Whether QoS needs to be deployed in the network depends on the applications and the business-critical nature of those applications. In order to deploy an effective QoS strategy within the Enterprise, it is imperative to understand the types of applications and the traffic patterns. For most installations, QoS is required for an IP Telephony deployment or any other application that is time/delay sensitive (e.g., video conferencing). QoS can also be employed for mission-critical applications such as Enterprise Resource Planning (ERP) tools. It is not necessary to provide QoS for every application on the network, only those that require special treatment. Note that even if there is sufficient bandwidth available in the network, QoS is still required for real-time applications to ensure minimum latency. It is also imperative to understand the potential impact of network component failures (uplinks, switches, modules) where an end-to-end

QoS deployment will ensure the high priority traffic (business applications) will be forwarded and best effort (or lower priority traffic) discarded in the event of congestion.

There are various strategies for deploying QoS throughout the infrastructure, and Avaya provides several tools to streamline the implementation. The Ethernet switches have the ability to either mark or honor the DSCP within each Ethernet packet. Many end devices now have the ability to set their own DSCP and thus set their own priority across the network. For example, the Avaya IP phones set their DSCP to Expedited Forwarding for all voice traffic, which maps into the Premium Service Class queue.

Care should be taken when simply honoring the markings from end stations. Current Windows operating systems can also mark DSCP, and therefore savvy users could prioritize their traffic on the network that honors the DSCP marking. It is a better practice for the edge switch to re-mark the DSCP to one that is controlled by the network administrator. This can be accomplished by using VLAN prioritization, which marks all packets in a specific VLAN with the same priority (prioritizing the voice VLAN), or by using the filtering capabilities of the Ethernet switches to mark packets individually based on filtering criteria established by the network administrator.

To assist in qualifying these types of applications and the associated QoS levels, Avaya has created a QoS matrix and has standardized Avaya Service Classes across all platforms. This matrix is intended as a guideline for the implementation of QoS:

Avaya Service Class	Target Applications and Services	Tolerance to:		
		Loss	Delay	Jitter
Critical	Super user Telnet, Critical heartbeats between routers/switches	Very Low	Very Low	Very Low
Network	ICMP, OSPF, BGP, RIP, ISIS, COPS, RSVP DNS, DHCP, BootP, high priority OAM	Low	Low	N/A
Premium	VoIP, T.38 Fax over IP, Lawful Intercept, CES Real-time VPN service (CIR > 0, EIR = 0)	Very Low	Very Low	Very Low
Platinum	Video Conferencing, Interactive Gaming Real-time VPN service (CIR > 0, EIR > 0)	Low	Low	Low
Gold	Streaming audio, video on demand Broadcast TV, video surveillance	Low-Med	Med-High	High
Silver	Credit card transactions, wire transfers Instant Messaging Low Loss/Delay Data VPN service (CIR > 0, EIR > 0)	Low	Low-Med	N/A
Bronze	E-mail Non-time-critical OAM&P	Low	Med-High	N/A
Standard	Best effort applications Best effort VPN (CIR >= 0, EIR > 0)	Med	High	N/A

Table 2.8: Quality of Service Matrix

The following highlights the IP header/DSCP and the DSCP/ToS/IP precedence mapping to the Avaya Service Classes. There are 64 possible different DSCP markings that can be utilized for QoS in the network, along with four different per hop behaviors:

- Expedited Forwarding – voice services
- Assured Forwarding – real-time and non-real-time applications
- Class Selector – used to support legacy routers
- Default Forwarding – best effort

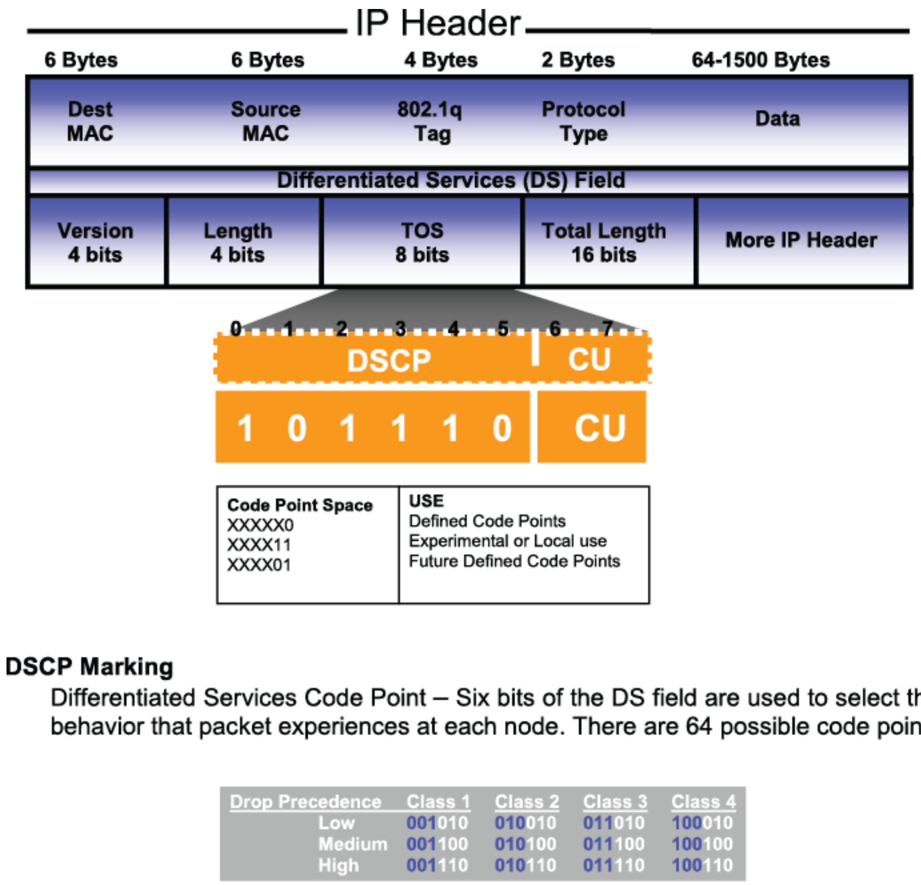

Code Point Space	USE
XXXXX0	Defined Code Points
XXXX11	Experimental or Local use
XXXX01	Future Defined Code Points

DSCP Marking

Differentiated Services Code Point – Six bits of the DS field are used to select the per hop behavior that packet experiences at each node. There are 64 possible code points.

Drop Precedence	Class 1	Class 2	Class 3	Class 4
Low	001010	010010	011010	100010
Medium	001100	010100	011100	100100
High	001110	010110	011110	100110

Figure 2.22: IP Header – DSCP Definition

Table 2.9 depicts the Avaya mapping of DSCP, Type of Service (ToS), and IP precedence to the Avaya Service Classes, along with their mapping into the DSCP per hop behavior.

DSCP	TOS	IP Precedence	Binary	NNSC
0x0	0x0	0	000000 **00**	Standard
0x0	0x0	-	000000 **00**	
0x8	0x20	1	001000 **00**	Bronze
0xA	0x28	-	001010 **00**	
0x10	0x40	2	010000 **00**	Silver
0x12	0x48	-	010010 **00**	
0x18	0x60	3	011000 **00**	Gold
0x1A	0x68	-	011010 **00**	
0x20	0x80	4	100000 **00**	Platinum
0x22	0x88	-	100010 **00**	
0x28	0xA0	5	101000 **00**	Premium
0x2E	0xB8	-	101110 **00**	
0x30	0xC0	6	110000 **00**	Network
0x38	0xE0	7	111000 **00**	Critical

DSCP and TOS are in HEX

IP Precedence in decimal

NNSC: Avaya Service Class

PHB: Per Hop Behavior

Table 2.9: Default Avaya DSCP / ToS / IP Mapping

The VSP 9000 classifies traffic as it enters the DiffServ network, and assigns appropriate Per-Hop Behavior (PHB), based on the classification. To differentiate between classes of service, the VSP 9000 marks the DiffServ (DS) parameter in the IP packet header, as defined in RFC2474 and RFC2475. The DSCP marking defines the forwarding treatment of the packet at each network hop. This marking (or classification) occurs at the edge of the DiffServ domain, and is based on the policy (or filter) associated with the particular microflow or aggregate flow.

You can configure the mapping of DSCP-to-forwarding behaviors and DSCP re-markings. Remarking the DSCP resets the treatment of packets based on new network specifications or desired levels of service. All IST/SMLT ports should be configured as DiffServ Trusted ports, which honor the DSCP marking in the packets.

Separate QoS filters can be created for any locally attached devices that require special treatment on the network. As shown in the figure below, the ports used to connect to the Server would be configured with QoS filters to mark specific DSCP values based on the application and/or traffic type.

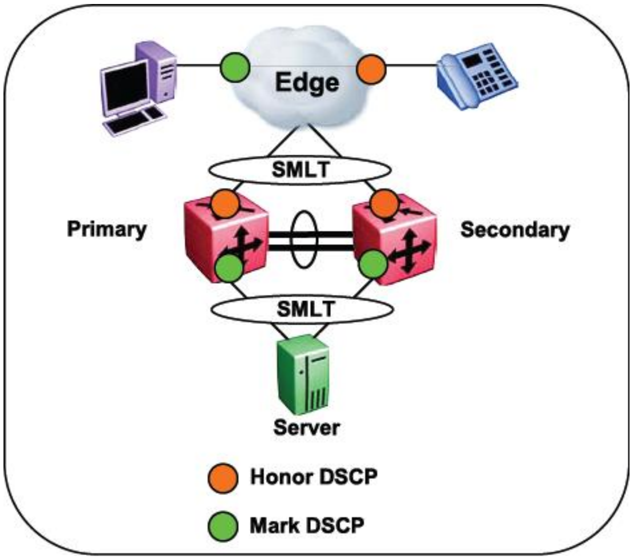

Figure 2.23: VSP 9000 Core QoS

The VSP 9000 has functions you can use to provide appropriate QoS levels to traffic for each customer, application, or packet. These functions include port-based shapers, DiffServ access or core port settings, policy-based policers, and port-based policers. The Avaya Virtual Services Platform 9000 also provides access control list (ACL)-based filters. You do not need to use filters to provide QoS; however, filters aid in prioritizing customer traffic. Filters also provide protection by blocking unwanted traffic.

Policers apply at ingress; shapers apply at egress. ACL-based filters apply at ingress and egress.

2.1.23 DHCP Relay

DHCP is used to provide IP addresses to end stations in the network. In the design scenario where multiple VLANs are used, DHCP relay will be required to get the DHCP information from the end user VLAN to the DHCP server (which would normally sit on a core server VLAN).

- Enable per VLAN with the physical VLAN IP address as the DHCP relay agent
- Multiple DHCP server addresses are supported (up to 10) per relay agent
- Do not use the VRRP virtual IP address as the DHCP Relay agent, always use the physical IP address of the VLAN

2.1.24 VRRP with Backup Master

VRRP provides redundancy for end user's default gateway and should be utilized for each VLAN configured that host end stations. Along with VRRP, Backup Master should be enabled on the Switch Cluster to provide active-active routing and forwarding of traffic.

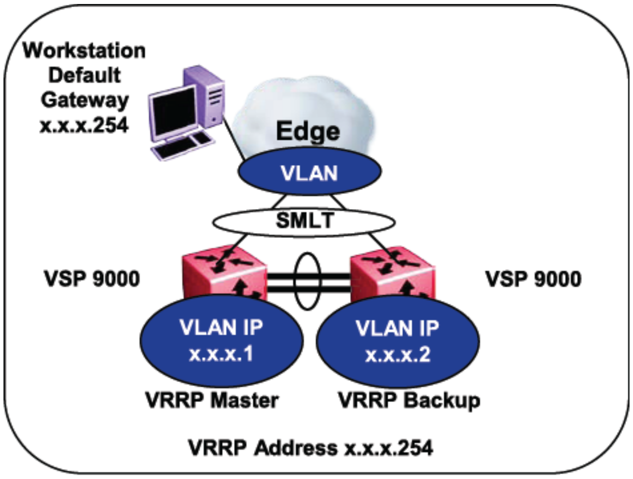

Figure 2.24: VSP 9000 VRRP

- Enable VRRP and Backup Master on each VLAN
- Configure VRRP priority higher than 100 (i.e. 200) to set VRRP Master
- Stagger VRRP Masters between VSP 9000's in the core
- Leave VRRP priority at default (100) for VRRP Backup
- Do not configure the virtual address as a physical interface that is used on any of the routing switches – use a third address, for example:
    - Physical IP address of VLAN on Switch 1 = x.x.x.1
    - Physical IP address of VLAN on Switch 2 = x.x.x.2
    - Virtual IP address of VLAN a = x.x.x.254

2.1.25 RSMLT Layer 2 Edge

RSMLT L2 Edge offers an alternative to VRRP for end user default gateway redundancy. VRRP and RSMLT L2 Edge can be used on the same Switch Cluster on different VLANs, but do not use both VRRP and RSMLT L2 Edge on the same VLAN simultaneously.

The RSMLT implementation does not use a Virtual IP address but instead uses physical IP addresses on each VSP 9000 for redundancy. RSMLT L2 Edge stores the RSMLT peer MAC/IP address-pair in its local configuration file and restores the configuration if the peer does not restore after a simultaneous reboot of both RSMLT peer switches. It is imperative to save the configuration file on each VSP 9000 when RSMLT L2 Edge if first implemented to ensure that the peer MAC/IP address-pair is saved in the configuration file.

Each VSP 9000 is able to forward on behalf of itself as well as its peer in the Switch Cluster. This makes very efficient use of bandwidth and resources and also ensures seamless failover and recovery in the event of a failure.

Avaya recommends using RSMLT L2 Edge in place of VRRP as it provides several advantages, including:

- RSMLT is only limited by the number of IP interfaces on the VSP 9000
- VRRP is limited to 250 instances
- RSMLT requires significantly less control traffic
- RSMLT is much less intensive on CPU resources

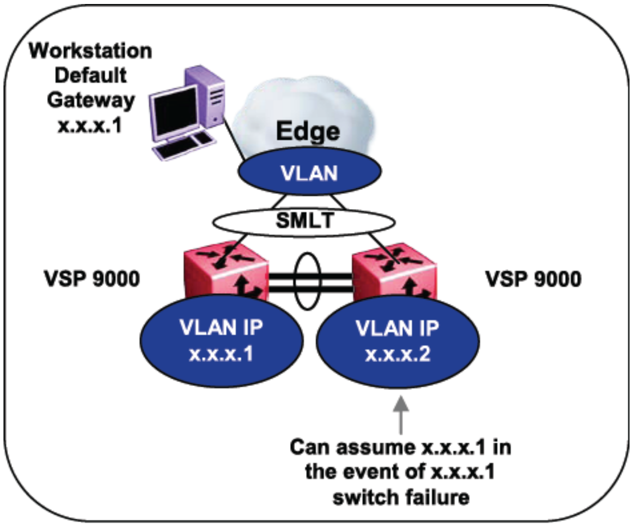

Figure 2.25: VSP 9000 RSMLT L2 Edge

When implementing RSMLT L2 Edge, make sure to:

- Enable RSMLT on each VLAN
- Enable RSMLT-Edge-Support globally
- Configure the Hold-up timer to 9999 (infinity) – this timer defines how long the RSMLT switch maintains forwarding for its peer

2.1.26 Layer 3 Routing

The Super Large Campus Design will most likely require a unicast routing protocol, the VSP 9000 presently supports OSPF, RIP, and BGP. OSPF is the Avaya recommended unicast routing protocol for the Super Large Campus design. The efficiencies of OSPF along with its scalability and overall market adoption make it the best alternative for routing.

Please consult the Release Notes for the VSP 9000 for route scalability, number of IP interfaces, number of OSPF areas, etc. supported per switch/stack.

When designing an OSPF network, review the following criteria:

- Timers must be consistent across the entire network – **use default values**
- Configure core switches as designated routers
- Higher rtrpriority is designated router
- Configure rtrpriority in increments of 10 for future flexibility
- Configure **OSPF router ID** the same as a 32-bit IP address assigned to a circuitless/loopback interface
- Must enable **ASBR** to provide static, RIP or BGP4 route redistribution
- Use **MD5 authentication** on any untrusted OSPF links
- Use **OSPF area summarization** to reduce routing table sizes
- Use **Stub or NSSA** areas as much as possible to reduce CPU overhead
- Use **OSPF passive interfaces** to ensure routing updates are not sent out VLANs where no routers are present
- Use **OSPF active interfaces** only on intended route paths. Typically, you should configure wiring closet subnets as OSPF passive interfaces unless they form a legitimate routing path for other routes.
- Limit the number of **OSPF areas per switch to as few as possible** to avoid excessive shortest path calculations. Be aware that the switch has to execute the Djikstra algorithm for each area separately.
- Ensure that the **OSPF dead interval** is at least four times the OSPF hello interval

When designing a RIP network, review the following criteria:

- For user VLANs where RIP is enabled, disable default supply and listen to reduce the amount of broadcast traffic on those VLANs
- Disable RIP supply and learn on all access ports
- Use triggered updates to help reduce routing convergence times

2.1.26.1 Equal Cost Multipath (ECMP) for Layer 3 link load balancing

ECMP enables the ability to load balance Layer 3 links and provide redundancy for routing in situations where there are at least two equal paths from the source to the destination network. ECMP uses a very similar algorithm as Multilink Trunking to distribute traffic among the links.

The number of paths is configurable. The VSP 9000 supports up to eight paths; the ERS 8300 and ERS 5000 support up to four equal cost paths. Note that ECMP is supported on the ERS 5000 series switches with the exception of the ERS 5510. ECMP can operate in a mixed ERS 5000 stack but cannot run on a ERS 5510. ECMP supports and complements OSFP, RIP, and Static routes.

2.1.26.2 Routed Split Multilink Trunking (RSMLT)

Avaya's RSMLT permits rapid failover for core topologies by providing an active-active router concept to core Split MultiLink Trunking (SMLT) networks. RSMLT scenarios include SMLT triangles, squares, and SMLT full-mesh topologies, with routing enabled on the core VLANs or edge VLANs when connected end device running an IGP. In the event of core router failures, RSMLT manages packet forwarding, thus eliminating dropped packets during the routing protocol convergence. RSMLT supports the following:

- IP Unicast Static Routes
- RIP v1/v2
- OSPF
- BGP

When RSMLT is enabled on a VLAN (on both aggregation devices) the cluster switches simply inform each other (over IST messaging) of their physical IP/MAC on that VLAN; thereafter the two Cluster switches take mutual ownership of each other's IP address on that VLAN; which means each cluster switch will:

- Reply to ARP request for both it's IP and it's Peer's IP on that VLAN
- Reply to pings to it's IP and it's Peer's IP on that VLAN
- Route IP traffic which is being directed to the physical MAC of it's IP or the physical MAC of it's peer's IP on that VLAN

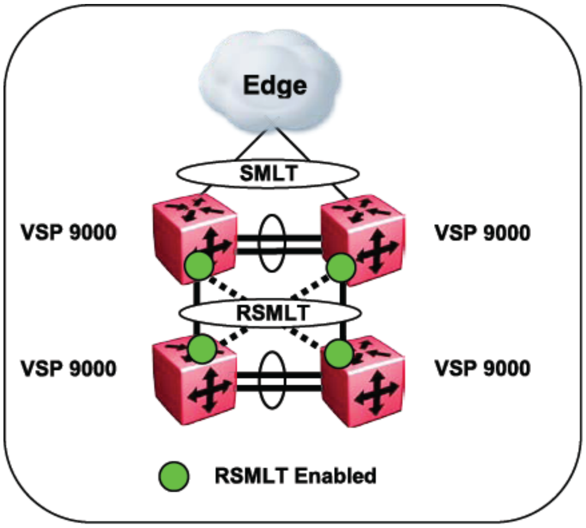

- End-to-end sub-second failover for routed traffic
- All switches in the core configured for normal routing (static, RIP, OSPF, BGP)
- No VRRP or ECMP required in the core
- Multiple RSMLT core VLANs supported per switch

Figure 2.26: VSP 9000 Routed Split Multilink Trunking (RSMLT)

RSMLT runs under the existing routing protocols, so no tuning of the IGP is necessary and there is no direct interaction between RSMLT and the IGP.

- Based on SMLT, so all SMLT rules apply
- Configured on a per VLAN basis
- VLAN must be routable and part of the SMLT links and IST link
- If possible, ensure that destination networks are directly accessible from each of the switches participating in RSMLT. This guarantees sub-second failover
- The hold-up timer defines how long the RSMLT-peer switch keeps routing for its peer after a peer switch failure. Leave the hold-up timer at default of 180 seconds when OSPF is used.
- The hold-down timer defines how long the switch waits before activating Layer 3 forwarding for its peers' MAC address. Leave the hold-downup timer at default of 60 seconds when OSPF is used. If RIP is used, the value should be changed to 180 seconds.
- In an RSMLT-based VLAN, the maximum number of supported SMLT/SLTs is 32. The RSMLT-based VLAN can have as many non-SMLT/SLT ports as needed.

2.1.26.3 RSMLT Dual Core VLANs

By using Avaya technologies in new innovative ways, the availability and resiliency of the network can be increased substantially. The following uses RSMLT in a configuration with dual core VLANs to minimize traffic interruption in the event of losing the OSPF designated router (DR).

When the DR fails, the backup designated router (BDR) is immediately elected as the new DR on the broadcast segment. As soon as the new DR is elected (via the Hello protocol) all routers must immediately perform a shortest path first (SPF) run and issue new link state advertisements (LSAs) for the segment. Since the DR is the only router generating the Network LSA for the segment, a new Network LSA needs to be generated by the new DR (the old one no longer being valid) and every router on that segment also needs to refresh their own router LSA.

The SPF run is done as soon as the router has detected a new DR. If the router is quick, it will perform the SPF run before receiving the new LSAs for the segment; as such all OSPF routes across this particular segment are not possible after the first SPF run. Whereas if the router is slow, by the time it detects the new DR and does it's SPF run, it already has received most (if not all) of the new LSAs from its neighbors and it's routing table will be still be able to route most (if not all) routes across the segment.

In either case, if the DR fails, two SPF runs are necessary to restore routes across the segment. But the OSPF hold-down timer will not allow two consecutive SPF runs to occur within the value of the timer. That timer defaults to 10secs on the VSP 9000 and can be reduced to a minimum of 3 seconds, however reducing this timer is not recommended. Every time the DR is lost, traffic interruption of up to 10 seconds can occur.

The solution for this scenario is to do the following:

- Create a second OSPF Core VLAN, forcing different nodes to become DR for each VLAN.
- Each OSPF Core VLAN will have DR (set priority to 100) and no BDRs (set OSPF priority to 0 on all routers/switches not intended to become the DR).
- No BDR is necessary – the two VLANs back each other up from a routing perspective.

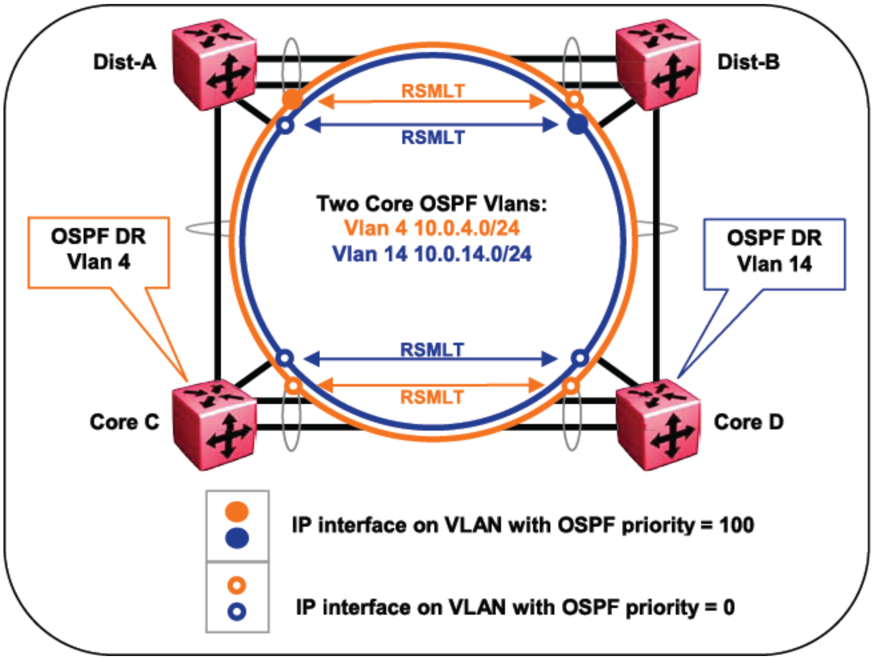

Figure 2.27: RSMLT with Dual Core VLANs

Dual core VLANs are not recommended when deploying multicast.

2.1.27 Multicast

Multicast routing is used to distribute multicast traffic within the network. Sources of multicast traffic must be routed on different IP subnets to reach the multicast subscribers. The VSP 9000 is uniquely designed to support multicast traffic flows in a very efficient manner. Ingressing packets are classified in hardware and only replicated where needed.

Avaya recommends using PIM as the multicast routing protocol of choice. PIM only sends multicast to areas of the network that have specifically requested the multicast stream. This is a much more efficient use of the available bandwidth in the network. PIM is usually the protocol of choice when there are a sparse number of users requesting the multicast stream. PIM-SSM is usually the protocol of choice for applications such as TV distribution or applications that require transmission acknowledgement. PIM uses the underlying routing table of the unicast routing protocol for its route table – for large scale networks, it is best to use OSPF.

The Super Large Campus Solution will likely have multicast application requirements. The VSP 9000 supports PIM-SM with Switch Clustering and IGMP over SMLT/SLT. The network design accommodates the multicast needs by using PIM-SM for multicast routing in the core and IGMP Snooping and Proxy at the edge. When implementing Multicast routing in the Super Large Campus, follow the recommendations outlined below:

- Use PIM-SM as the Layer 3 Multicast routing protocol
- Enable IGMP Snooping and Proxy on the edge switches

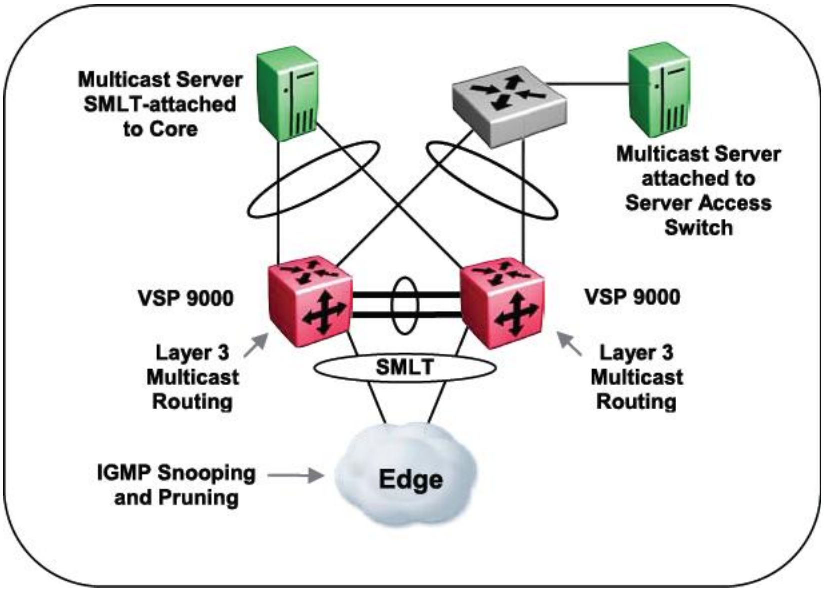
Multicast Server SMLT-attached to Core
Multicast Server attached to Server Access Switch
VSP 9000
VSP 9000
Layer 3 Multicast Routing
Layer 3 Multicast Routing
SMLT
IGMP Snooping and Pruning
Edge

Figure 2.28: VSP 9000 Multicast Routing

2.1.28 Server Connectivity

The preferred method for connecting servers in the Super Large Campus Solution is to use a subtended Server Access Switch/Stack (ERS 2500/4500/5000). For horizontal stacking information and configuration details, see the *Data Center Server Access Solution Guide* (NN48500-577). Servers can also be directly attached to the core if required. However, keeping the core functionality separated from server connectivity is preferred when possible to aide in fault isolation, maintenance, and troubleshooting.

Server virtualization is also fully supported with the Super Large Campus Solution. Avaya and VMware have fully certified a solution as described in the Resilient Data Center Server Edge Solutions for VMware ESX Server Technical Configuration Guide (NN48500-542).

General design considerations:

- ➢ Disable Spanning Tree on the ports connected to the servers (MLT/SMLT/SLT)
- ➢ Server NIC Teaming configurations supported include:
  - Broadcom NICs using FEC/GEC Generic Trunking
  - Intel NICs using Static Link Aggregation
  - HP NICs using SLB with automatic Transmit Load Balancing

Quantity of Servers supported when connecting to a subtended Server Access Switch/Stack:

- ➢ MLT-based servers will be the number of MLT groups supported on the Server Access Switch/Stack less one needed for the SMLT to the Switch Cluster Core.
  - ERS 5000 supports 32 MLT groups with a maximum of 8 ports per group
  - ERS 4500 supports 32 MLT groups with a maximum of 8 ports per group
  - ERS 2500 supports 6 MLT groups with a maximum of 4 ports per group

- Port-based servers will be the number of SLT groups supported on the Server Access Switch/Stack less two required for IST. Every port (saving at least two for the IST) can be used for SLT groups terminating into the cluster, with each SLT group consisting of a maximum of two uplinks (one per core Ethernet Routing Switch).

  The ERS 5000 supports up to 398 Copper port-based SLT groups, 190 1GbE port-based SLT groups, and 62 10GbE port-based SLT groups.

Quantity of Servers supported when connecting directly to Switch Cluster Core:

- SMLT-attached servers equals the number of MLT groups supported on the core less the number of MLT groups already in use by Edge switch connections via SMLT or MLT.
- SLT-attached servers will be number of ports available on the Switch Cluster Core

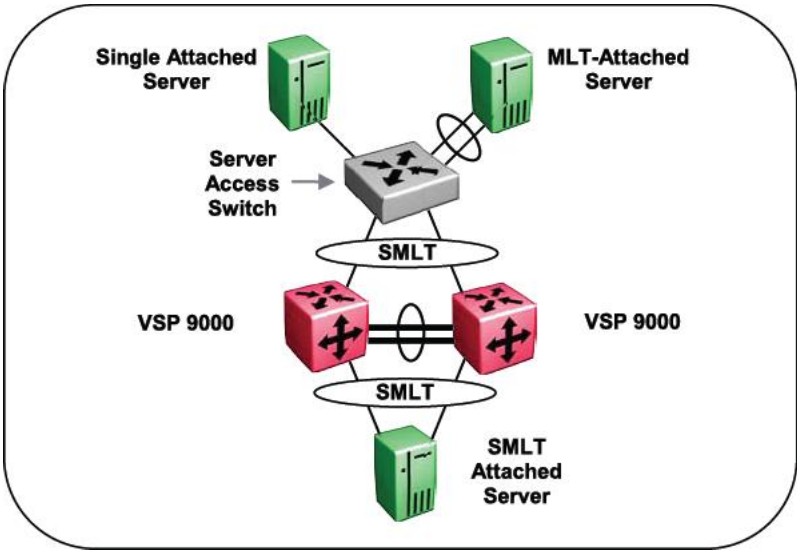

Figure 2.29: VSP 9000 Server Connectivity

2.2 Edge Switching

The Edge Switching provides the end user connectivity for the Super Large Campus Solution. Avaya provides flexibility in which products you can use at the edge: ERS 2500, ERS 4500, ERS 5000 Series, and the ERS 8300.

Providing redundancy in the hardware is the basic building block to creating the highly resilient core and therefore the platforms deployed must provide redundancy in hardware components. These redundancy features will vary between chassis-based and stackable/standalone products. This does not imply that one platform is more or less redundant than any of the other platforms; it simply states that the way redundancy is achieved will vary based on the hardware being used.

The ability to hot-swap any and all hardware components is an absolute necessity and should be one of the highest criteria when evaluating hardware platforms.

For chassis systems, key hardware redundancy components include; power supplies, fan trays, switch fabric/CPUs, and I/O modules. For stackable/standalone systems, the key hardware redundancy components include; external/internal redundant power supplies, resilient stacking architecture and the ability to hot-swap any switch without interrupting traffic flow from/to other switches in the stack. Avaya is the first vendor to break through the terabit stackable boundary.

The Edge switching portfolios support a multitude of features, all of which are not presented in this solution. The features and functionality highlighted here represent the basic requirements for the Super Large Campus Solution. Please refer to the product documentation for a detailed explanation of these and all the other features of the above mentioned platforms.

This section will detail the switching options for the edge and also provide the best practice recommendations. Platform differentiations and feature differences will be highlighted as applicable in each section for Edge Switching.

- Edge Switching Products
- Stacking of Edge Switches
- Virtual LANs (VLANs)
- Link Aggregation
- Spanning Tree
- BPDU Filtering
- VLANs
- DHCP Relay
- Quality of Service
- Layer 3
- VRRP with Backup Master

2.2.1 Edge Switching Products

The ERS 2500 product portfolio offers the following four different switch models:

ERS 2526T

- 24 ports 10/100
- 2 combo ports (SFP or 10/100/1000)

ERS 2526T-PWR

- 24 ports 10/100 (12 with PoE)
- 2 combo ports (SFP or 10/100/1000)

ERS 2550T

- 48 ports 10/100
- 2 combo ports (SFP or 10/100/1000)

ERS 2550T-PWR

- 48 ports 10/100 (24 with PoE)
- 2 combo ports (SFP or 10/100/1000)

The ERS 4500 product portfolio offers the following eleven different switch models:

ERS 4526T

- 24 ports 10/100
- 2 combo ports (SFP or 10/100/1000)

ERS 4526T-PWR

- 24 ports 10/100 with PoE
- 2 combo ports (SFP or 10/100/1000)

ERS 4550T

- 48 ports 10/100
- 2 combo ports (SFP or 10/100/1000)

ERS 4550T-PWR

- 48 ports 10/100 with PoE
- 2 combo ports (SFP or 10/100/1000)

ERS 4524GT

- 24 ports 10/100/1000
- 4 shared SFP ports

ERS 4524GT-PWR

- ➢ 24 ports 10/100/1000 with PoE
- ➢ 4 shared SFP ports

ERS 4548GT

- ➢ 48 ports 10/100/1000
- ➢ 4 shared SFP ports

ERS 4548GT-PWR

- ➢ 48 ports 10/100/1000 with PoE
- ➢ 4 shared SFP ports

ERS 4526GTX

- ➢ 24 ports 10/100/1000
- ➢ 2 XFP ports (10 Gig LAN Phy)

ERS 4526GTX-PWR

- ➢ 24 ports 10/100/1000 with PoE
- ➢ 2 XFP ports (10 Gig LAN Phy)

ERS 4526FX

- ➢ 24 ports 100FX
- ➢ 2 combo ports (SFP or 10/100/1000)

The ERS 5000 product portfolio offers the following ten different switch models:

ERS 5650TD

- ➢ 48 ports 10/100/1000
- ➢ 2 XFP ports (10Gig)

ERS 5650TD-PWR

- ➢ 48 ports 10/100/1000 with PoE
- ➢ 2 XFP ports (10Gig)

ERS 5698TFD

- ➢ 96 ports 10/100/1000
- ➢ 6 Shared SFP ports
- ➢ 2 XFP ports (10Gig)

ERS 5698TFD-PWR

- 96 ports 10/100/1000 with PoE
- 6 Shared SFP ports
- 2 XFP ports (10Gig)

ERS 5632FD

- 24 SFP ports
- 8 XFP ports (10Gig)

ERS 5530-24TFD

- 12 10/100/1000 ports
- 12 Shared (SFP or 10/100/1000)
- 2 XFP ports (10Gig)

ERS 5520-48T-PWR

- 48 10/100/1000 ports with PoE
- 4 Shared SFP ports

ERS 5520-24T-PWR

- 24 10/100/1000 ports with PoE
- 4 Shared SFP ports

ERS 5510-48T

- 48 10/100/1000 ports
- 2 Shared SFP ports

ERS 5510-24T

- 24 10/100/1000 ports
- 2 Shared SPF ports

The ERS 8300 product portfolio offers the following eleven different switch models:

ERS 8300

- 6 slot or 10 slot chassis
- Dual switching fabrics with integrated uplink ports for maximum density
- N+1 power supply redundancy (AC and DC options available)
- All components fully hot swappable

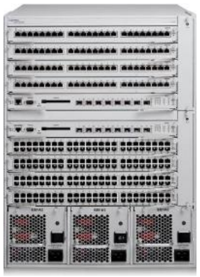

Module	Ports	Type
8393SF	8	288Gbps Switch Fabric with 8 1GbE SFP ports
8394SF	2	288Gbps Switch Fabric with 2 10GbE XFP ports
8348TX	48	48 port 10/100BaseT
8348TX-PWR	48	48 port 10/100BaseT with 802.3af PoE
8324GTX	24	24 port 10/100/1000BaseT
8348GTX	48	48 port 10/100/1000BaseT
8348GTX-PWR	48	48 port 10/100/1000BaseT with 802.3af PoE
8348GB	48	48 port 1GbE SFP
8308XL	8	8 port 10GbE XFP

8308XL

8394SF

8393SF

8348GTX-PWR

2.2.2 Stacking of Edge Switches

The ERS 2500, ERS 4500, and ERS 5000 switches support a common resilient stacking architecture. The stack is created by using the stacking cables and stacking ports on the ERS switches. Switches are cabled together in the manner shown below so that every switch has two stacking connections for utmost resiliency. The shortest path algorithm used for stacking allows for the most efficient use of bandwidth across the stack. The maximum number of switch can be allowed to stack is 8 and can be of any model (mix and match) within the same product family.

A failure in any unit of the stack will not adversely affect the operation of the remaining units in the stack. Replacement of the failed switch is easy with the Auto-Unit Replacement feature. This allows for a new switch to be put into the stack and will automatically get the right software image and configuration without user intervention – the replacement switch must be the exact model of the failed switch. In addition, the SW image has to be right for the AUR to work properly. Please refer to Product documentation for more info on AUR.

- Enable Stack Monitor on all Edge stacks which will alert on any stack size changes
- ERS 2500
  - 4Gbps stacking bandwidth per switch, 32Gbps maximum per stack
  - Requires a stack license per unit
- ERS 4500
  - 40Gbps stacking bandwidth per switch, 320Gbps maximum per stack
  - Enable Forced Stack Mode when using stacks of two switches

- ERS 5000
  - ERS 5500 – 80Gbps stacking bandwidth per switch, 640Gbps maximum per stack
  - ERS 5600 – 144Gbps stacking bandwidth per switch, 1.1Tbps maximum per stack
  - Hybrid stacks of 5500 and 5600 switches are supported
  - Enable Forced Stack Mode when using stacks of two switches

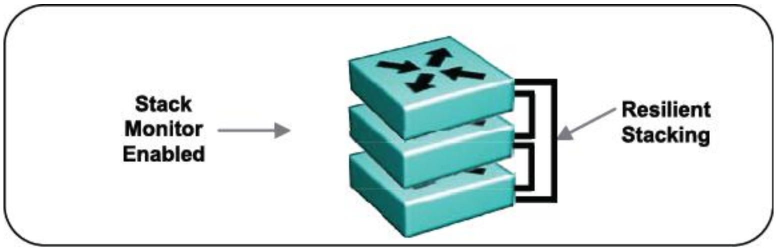

Figure 2.30: Edge Stacking

2.2.3 Power over Ethernet

The use of Power over Ethernet (PoE) has become increasingly popular over the past few years and is standardized by the IEEE as 802.3af. Many end devices, such as wireless access points, security cameras, and security card readers now support Power over Ethernet; however, the largest market driver for PoE today is IP Telephony. The ability to centrally power IP phones from the Ethernet switch has made PoE ubiquitous in IP Telephony environments.

There are two components to PoE:

- Power Sourcing Equipment (PSE) which is normally the Ethernet switch providing the PoE
- Powered Device (PD) which is the end device using the power

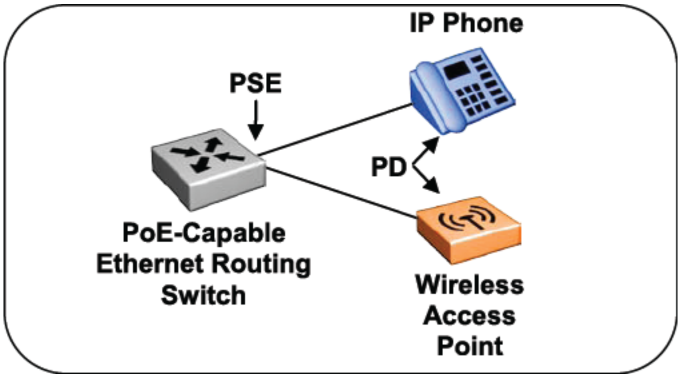

Figure 2.31: Power over Ethernet

The PSE will not send power out of the Ethernet switch port until it is able to verify that the end station is PoE capable. The 802.3af standard specifies a resistive discovery mechanism for the PSE to perform this function. This resistive method sends out two low voltage electrical pulses and once the PSE has successfully discovered the PD, it will begin providing power. Legacy (pre-standard) PoE devices use a capacitive detection mechanism to detect the PD. All Avaya PoE switches support both resistive and capacitive detection mechanisms as many pre-standard PDs supported only capacitive discovery.

Power can be provided over the used data pairs (1, 2, 3, and 6) or the unused pairs (4, 5, 7, and 8) in a UTP copper cable. The 802.3af standard mandates that PDs must be able to accept power from either option. The Avaya Ethernet switches (2500, 4500, 5000, and 8300) provide power over the used pairs.

When designing a Converged Campus network, it is imperative to understand the PoE requirements in order to provide adequate power to the end devices. The 802.3af standard specifies a maximum of 15.4 watts per device depending on power classification, with many devices using only 4 to 8 watts.

The 802.3af standard defines the different classes of detection, as defined in Table 2.10. The Avaya Ethernet switches do not use power classification to provide PoE to the PDs. The Ethernet switches use a pool of power per switch or per module. As devices come online and begin to use power, the overall pool of power is decremented. Power management and power priority are used to control the amount of power and what ports have priority to that power.

The PSE will not send power out of the Ethernet switch port until it is able to verify that the end station is PoE capable. The 802.3af standard specifies a resistive discovery mechanism for the PSE to perform this function. This resistive method sends out two low voltage electrical pulses and once the PSE has successfully discovered the PD, it will begin providing power. Legacy (pre-standard) PoE devices use a capacitive detection mechanism to detect the PD. All Avaya PoE switches support both resistive and capacitive detection mechanisms as many pre-standard PDs supported only capacitive discovery.

Power can be provided over the used data pairs (1, 2, 3, and 6) or the unused pairs (4, 5, 7, and 8) in a UTP copper cable. The 802.3af standard mandates that PDs must be able to accept power from either option. The Avaya Ethernet switches (2500, 4500, 5000, and 8300) provide power over the used pairs.

When designing a Converged Campus network, it is imperative to understand the PoE requirements in order to provide adequate power to the end devices. The 802.3af standard specifies a maximum of 15.4 watts per device depending on power classification, with many devices using only 4 to 8 watts.

The 802.3af standard defines the different classes of detection, as defined in Table 2.10. The Avaya Ethernet switches do not use power classification to provide PoE to the PDs. The Ethernet switches use a pool of power per switch or per module. As devices come online and begin to use power, the overall pool of power is decremented. Power management and power priority are used to control the amount of power and what ports have priority to that power.

Class	Usage	Maximum power level with 100 m cable of Cat 5 at:	
		Output of PSE	Input of PD
0	Default	15.4 Watts	0.44–12.95 Watts
1	Optional	4.0 Watts	0.44–3.84 Watts
2	Optional	7.0 Watts	3.84–6.49 Watts
3	Optional	15.4 Watts	6.49–12.95 Watts

Table 2.10: PoE Classes of Power Input/output

- Four classes of DC detection (Classes 0, 1, 2 and 3) per IEEE 802.3af
- Four classes of AC detection (Classes 0, 1, 2 and 3) per IEEE 802.3af
- One class of Capacitive detection (Class 2) for pre-IEEE Avaya IP phones

The difference in output power of the PSE and input power of the PD is to account for loss in the Ethernet cable.

When designing the network, take into consideration the potential number of end devices requiring PoE and the amount of power each of these devices requires. This will help decide the number of PoE-capable ports required in the wiring closet.

The network administrator can also use features within the Ethernet switches to control the PoE across the switch. Power Management will enable/disable or limit the PoE on a per port/module basis. Power Priority allows certain ports to have priority over other ports when requiring PoE. In certain configurations, there may not be enough PoE available. The power priority feature will allow designated ports to get power over lower priority ports. This will ensure the critical PoE devices will always have power.

Another critical aspect of PoE is redundant power in the wiring closet. When relying on the Ethernet switch to provide power to the end devices, it becomes even more important to have redundant power available in the closet in case of an electrical circuit failure. The addition of redundant power can also be used by the Ethernet switches to add overall power capability and increase the total amount of power available for end devices using PoE.

In all cases, Avaya strongly recommends having separate electrical circuits available in the closet for the various AC feeds into the Ethernet switching equipment as well as the Redundant Power Supply Units (if applicable).

Avaya created a PoE Calculator tool which assists during the design phase. By simply selecting the number of Avaya IP phones and wireless access points, or by entering in information for third party PoE devices, the calculator will provide information regarding the power required, number of switches, and if necessary, the amount of redundant power required. This tool eliminates the guess work around the design of the PoE infrastructure. The PoE Calculator can be found on the support site of Avaya.com or by searching for document number NN48500-520.

A new PoE standard is now emerging. 802.3at, or also commonly referred to as PoE+ will increase the amount of power supplied by the PSE. Avaya will offer products supporting 802.3at in the near future. Please note that this will require new Ethernet switching hardware to support 802.3at and the increased power. More information on this emerging standard can be found at http://www.ieee802.org/3/at/

The design of a network capable of PoE will vary depending on the Ethernet switching equipment you choose. The following sections highlight the various Avaya PoE options available:

Ethernet Routing Switch 8300

This chassis system provides both 10/100 and 10/100/1000 48 port I/O modules capable of PoE. When utilizing PoE, make sure to engineer the power requirements of the chassis properly. The amount of PoE per module is configurable up to 800 watts per module, along with the ability to specify port priority for PoE. The total PoE power required will dictate the type of input power for the chassis. The ERS 8300 provides different power options as indicated in Table 2.11.

ERS 8300 Six Slot Chassis

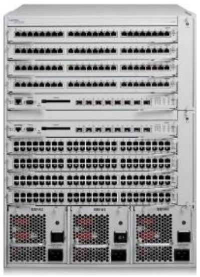

ERS 8300 Ten Slot Chassis

Power Supply	Power Supply Rating	# of Power Supplies	Redundancy	PoE Available
8301AC	110-120 VAC 20 Amp 1140 watts	1	No	400 watts
		2	Yes 1+1	400 watts
		3	Yes 2+1	800 watts
	200-240 VAC 20 Amp 1770 watts	1	No	800 watts
		2	Yes 1+1	800 watts
		3	Yes 2+1	1600 watts
8302AC	100-120 VAC 15 Amp 850 watts	1	No	200 watts
		2	Yes 1+1	200 watts
		3	Yes 2+1	400 watts
	200-240 VAC 15 Amp 1400 watts	1	No	400 watts
		2	Yes 1+1	400 watts
		3	Yes 2+1	800 watts

Table 2.11: ERS 8300 Power over Ethernet Options

Ethernet Routing Switch 5600

The PoE capable ERS 5600 series stackable switches are available in a 48-port and a 96-port version. The ERS 5600 offers built-in, hot swappable redundant power supply options in both AC and DC varieties. It is also capable of providing full 15.4watts per port on every port in the switch along with full N+1 redundant power simultaneously. The available configurations for power options are specified in Table 2.12.

ERS 5650TD-PWR

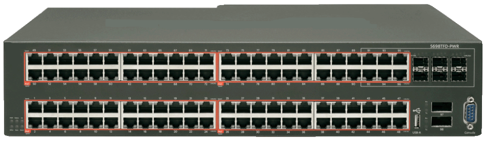

ERS 5698TFD-PWR

Switch Model	PoE with one power supply	PoE with two power supplies	PoE with three power supplies
ERS 5650TD-PWR (600W)	370 watts total 7.7 watts/port	740 watts total 15.4 watts/port	N/A
ERS 5650TD-PWR (1000W)	740 watts total 15.4 watts/port	740 watts total * 15.4 watts/port	N/A
ERS 5698TFD-PWR (1000W)	740 watts total 7.7 watts/port	1480 watts total 15.4 watts/port	1480 watts total * 15.4 watts/port
* Full 15.4 watts on every port with N+1 power redundancy			

Table 2.12: ERS 5600 Power over Ethernet Options

Ethernet Routing Switch 5500

The PoE capable ERS 5520 stackable switch is available in both a 24-port and a 48-port version. The ERS 5520 provides up to 320 watts per switch on standard 110 VAC power. To provide more power and/or redundant power, use the ERS Redundant Power Supply 15 (RPS 15) to augment the ERS 5520. The RPS 15 can support up to three ERS 5520 switches. The available configurations for power options are specified in Table 2.13.

ERS 5520-24T-PWR

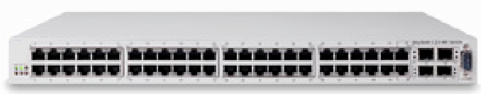

ERS 5520-48T-PWR

Switch Model	PoE on Standard AC	RPS 15 Power Sharing	RPS 15 RPSU
ERS 5520-24T-PWR	320 watts total 13.3 watts/port	740 watts total 15.4 watts/port	320 watts total 13.3 watts/port
ERS 5520-48T-PWR	320 watts total 6.7 watts/port	740 watts total 15.4 watts/port	320 watts total 6.7 watts/port

Table 2.13: ERS 5500 Power over Ethernet Options

Ethernet Routing Switch 4500

The PoE capable ERS 4500 stackable switches are available in 10/100 and 10/100/1000 48-port versions. The ERS 4500 provides up to 370 watts per switch on standard 110 VAC power. To provide more power and/or redundant power, use the ERS Redundant Power Supply 15 (RPS 15) to augment the ERS 4500. The RPS 15 can support up to three ERS 4500 switches. The available configurations for power options are specified in Table 2.14.

ERS 4526T-PWR

ERS 4550T-PWR

ERS 4524GT-PWR

ERS 4548GT-PWR

ERS 4526GTX-PWR

Switch Model	PoE on Standard AC	RPS 15 Power Sharing	RPS 15 RPSU
ERS 4526T-PWR	370 watts total 15.4 watts/port	740 watts total 15.4 watts/port	370 watts total 15.4 watts/port
ERS 4550T-PWR	370 watts total 7.7 watts/port	740 watts total 15.4 watts/port	370 watts total 7.7 watts/port
ERS 4524GT-PWR	360 watts total 15.0 watts/port	740 watts total 15.4 watts/port	360 watts total 15.0 watts/port
ERS 4548GT-PWR	320 watts total 6.7 watts/port	740 watts total 15.4 watts/port	320 watts total 6.7 watts/port
ERS 4526GTX-PWR	360 watts total 15.0 watts/port	740 watts total 15.4 watts/port	360 watts total 15.0 watts/port

Table 2.14: ERS 4500 Power over Ethernet Options

Ethernet Routing Switch 2500

The PoE capable ERS 2500 switches are available in both a 24-port and a 48-port version. With both of these ERS 2500 switches, PoE is provided on half the ports (ports 1-12 of the 24 port switch and ports 1-24 on the 48 port switch). The ERS 2500 provides up to 165 watts per switch on standard 110 VAC power. The ERS 2500 does not support a redundant power option. The available configurations for power options are specified in Table 2.15.

ERS 2526T-PWR

ERS 2550T-PWR

Switch Model	PoE on Standard AC	RPS 15 Power Sharing	RPS 15 RPSU
ERS 2526T-PWR	165 watts	N/A	N/A
ERS 2550T-PWR	165 watts	N/A	N/A

Table 2.15: ERS 2500 Power over Ethernet Options

Redundant Power Supply 15 (RPS 15)

The RPS 15 provides redundant power to the Avaya stackable Ethernet switches (both PoE and non-Poe). The RPS 15 is comprised of the following components:

- RPS 15 Chassis (supports up to three 600 watt power supplies)
- 600 Watt Power Supply
- DC-DC Converter (only required for some switches – see table below)
- DC cable to connect power supply to Ethernet switch

The RPS 15 supports two different DC cable types. The first (AA0005018) is used with all Ethernet switches that have a built-in DC-DC converter and can provide a single power connection to one Ethernet switch. The second type of cable, which comes in two models (AA0005020 – 25’ and AA0005021 – 10’) is used with all Ethernet switches that require the addition of the DC-DC converter module. This second cable type can provide a single power connection for up to four Ethernet switches.

The RPS 15 can be added to an Ethernet switch or stack of Ethernet switches while the switches are powered up and running. There is no need to power off the switch to connect the RPS 15 cable.

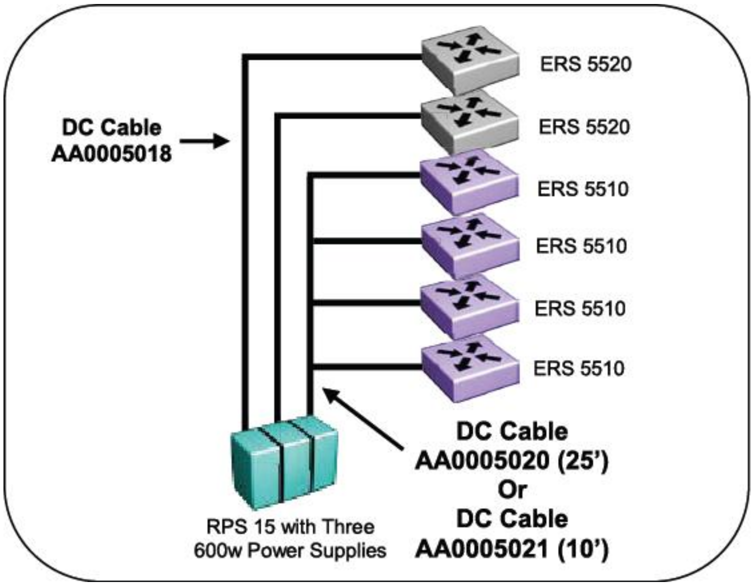
ERS 5520
ERS 5520
ERS 5510
ERS 5510
ERS 5510
ERS 5510
DC Cable AA0005018
DC Cable AA0005020 (25') Or DC Cable AA0005021 (10')
RPS 15 with Three 600w Power Supplies

Figure 2.32: Redundant Power Supply 15 (RPS15)

Table 2.16 provides information on the required components when using the RPS 15 with the various Ethernet switching options:

Switch Model	PoE Capable Switch	RPS 15 Chassis	RPS 15 600w Power Supply	DC-DC Converter	DC Cable for Built-In Converter	10' or 25' DC Cable
ERS 5510	No	1	1 per 4 switches	Required	N/A	Required
ERS 5520	Yes	1	1	Built-In	Required	N/A
ERS 5530	No	1	1	Built-In	Required	N/A
ERS 4526FX	No	1	1 per 4 switches	Required	N/A	Required
ERS 4526T	No	1	1 per 4 switches	Required	N/A	Required
ERS 4526T-PWR	Yes	1	1	Built-In	Required	N/A
ERS 4550T	No	1	1 per 4 switches	Required	N/A	Required
ERS 4550T-PWR	Yes	1	1	Built-In	Required	N/A
ERS 4524GT	No	1	1 per 4 switches	Required	N/A	Required
ERS 4524GT-PWR	Yes	1	1	Built-In	Required	N/A
ERS 4548GT	No	1	1 per 4 switches	Required	N/A	Required
ERS 4548GT-PWR	Yes	1	1	Built-In	Required	N/A

ERS 4526GTX	No	1	1 per 4 switches	Required	N/A	Required
ERS 4526GTX-PWR	Yes	1	1	Built-In	Required	N/A
ES 470-24T-PWR	Yes	1	1	Built-In	Required	N/A
ES 470-48T-PWR	Yes	1	1	Built-In	Required	N/A
ES 470-24T	No	1	1 per 4 switches	Required	N/A	Required
ES 470-48T	No	1	1 per 4 switches	Required	N/A	Required

Table 2.16: RPS 15 Configuration Options

The following table highlights the PoE requirements of several Avaya end devices:

PoE Device	Average PSE
1100 Series Phones	
IP Phone 1110 PEC NTYS02xxE6	2.80 watts
IP Phone 1120E PEC NTYS03xA -- NTYS03xCE6	7.00 watts
IP Phone 1120E PEC NTYS03xDE6 (see Note 7)	4.60 watts
IP Phone 1120E PEC NTYS03xEE6 (see Note 8)	4.20 watts
IP Phone 1140E PEC NTYS05xA -- NTYS05xCE6	7.30 watts
IP Phone 1140E PEC NTYS05xCE6 Rel 50 & higher (see Note 7)	4.80 watts
IP Phone 1140E PEC NTYS05xEE6 (see Note 8)	4.30 watts
IP Phone 1150E PEC NTYS06xxE6	7.00 watts
IP Phone 1165E PEC NTYS07xxE6 (see Note 9)	3.80 watts
1200 Series Phones	
1210 phone w/ integrated 3 port 10/100 switch	3.40 watts
1220 phone w/ integrated 3 port 10/100 switch	3.40 watts
1230 phone w/ integrated 3 port 10/100 switch	3.40 watts
1600 Series Phones	
Avaya IP Phone 1603 Class 2 w/ Adapter	4.32 watts
Avaya IP Phone 1608 Class 2	4.66 watts
Avaya IP Phone 1616 Class 3 (see Note 1)	6.22 watts
Avaya IP Phone 1616 w/ BM32 Class 3	6.82 watts
Avaya IP Phone 1616 Class 2 (see Note 1)	3.17 watts
Avaya IP Phone 1616 w/ BM32 Class 2	4.37 watts
Avaya IP Phone 1692 Class 3	5.20 watts
Gig Adapter for 1600 Series Phones	3.16 watts

<u>4600/5600 Series Phones</u>	
Avaya IP Phone 4601, 4602, 5601, 5602 Class 2	3.50 watts
Avaya IP Phone 4602SW, 5602SW Class 2	4.10 watts
Avaya IP Phone 4606, 4612, 4624 Class 0	5.00 watts
Avaya IP Phone 4610SW, 5610SW Class 2	4.00 watts
Avaya IP Phone 4620 Class 3	7.70 watts
Avaya IP Phone 4620SW Class 3 (see Note 2)	5.90 watts
Avaya IP Phone 4620SW, 5620SW Class 2	4.60 watts
Avaya IP Phone 4621SW, 4622SW Class 2	4.90 watts
Avaya IP Phone 4625SW Class 3	7.80 watts
Avaya IP Phone 4630SW Class 3	11.80 watts
Gig Adapter for 1600 Series Phones	3.16 watts
<u>9600 Series Phones</u>	
Avaya IP Phone 9610 Class 2	3.72 watts
Avaya IP Phone 9620 Class 2	4.90 watts
Avaya IP Phone 9620L Class 1	2.20 watts
Avaya IP Phone 9620C Class 2	4.60 watts
Avaya IP Phone 9630 Class 2	5.20 watts
Avaya IP Phone 9630G Class 2	4.83 watts
Avaya IP Phone 9640 Class 2	4.52 watts
Avaya IP Phone 9640G Class 2	4.83 watts
Avaya IP Phone 9650 Class 2	4.60 watts
Avaya IP Phone 9650C Class 2	4.52 watts
Avaya IP Phone 9670G Class 2	6.24 watts
<u>Phase 0 Phones</u>	
IP Phone 2004	3.30 watts
IP Phone 2004 w/ external 3 port 10/100 switch	11.00 watts
<u>Phase I Phones</u>	
IP Phone 2004 w/ integrated 3 port 10/100 switch	3.30 watts
IP Phone 2002 w/ integrated 3 port 10/100 switch	3.30 watts
<u>Phase II Phones</u>	
IP Phone 2001 w/ integrated 3 port 10/100 switch	3.20 watts
IP Phone 2002 w/ integrated 3 port 10/100 switch	3.20 watts
IP Phone 2004 w/ integrated 3 port 10/100 switch	3.40 watts
IP Phone 2007 w/ integrated 3 port 10/100 switch	7.00 watts

2000 Series Audio Conference Phone	
IP Audio Conference Phone 2033	11.6 watts
LG-Avaya Phones	
LG-Avaya IP Phone 6804	2.30 watts
LG-Avaya IP Phone 6812	2.30 watts
LG-Avaya IP Phone 6830	2.30 watts
LG-Avaya IP Phone 8540	4.50 watts
Wireless LAN	
WLAN 8120 AP	9.30 watts
WLAN 2220 AP	8.50 watts
WLAN 2230 AP	10.00 watts
WLAN 2330 – 2330A - 2332 AP	8.00 watts

Table 2.17: PoE Consumption for Avaya IP Phones and Access Points

2.2.4 Physical Layer Considerations/Fiber Fault Detection

Avaya provides several options for uplink connectivity over fiber – this does not necessarily preclude the use of copper for uplink connections; however, due to the distance limitations of copper (100 m), fiber is normally the media of choice. Both ends of a link generally must use the same transceiver type (that is, SX to SX), but they do not necessarily have to be from the same manufacturer or vendor – Avaya GBICs, SFPs, XFPs interoperate with most third-party vendors' products of the same transceiver type.

Avaya also supports the interoperability of CWDM with both XD and ZX GBICs. More specifically, one end of the link can use a ZX GBIC and the other end of the link a CWDM SFP, or one end can use an XD GBIC and the other end of the link a CWDM SFP. The following rules apply:

- XD GBIC with 40 Km CWDM SFP
- ZX GBIC with 70 Km CWDM SFP

Avaya offers SFPs and XFPs with a Digital Diagnostic Interface (DDI). These SFPs can provide significant diagnostic information when enabled by the Ethernet Switches via Digital Diagnostic Monitoring (DDM). At this time, the VSP 9000 supports DDM and allows the switch to monitor SFP and XFP laser operating characteristics. VSP 9000 support for Digital Diagnostic Interfaces (DDI—an interface that supports DDM) involves data collection and alarm and warning monitoring. Static data collection includes SFP vendor information, DDI support information, and DDI alarm and warning threshold values. Dynamic data collection includes temperature, supply voltage, laser bias current, transmit power, and receive power. DDM works at any time during active laser operation without affecting data traffic. The switch only checks warning and alarm status bits during initialization and during requests for dynamic data. If an alarm or warning is asserted or cleared, the switch logs a message and generates a trap. The switch maps DDM warning and alarm messages into Warning and Fatal message categories for system logging purposes. If you activate the ddm-alarm-portdown option, DDI shuts down the corresponding port if a high or low alarm occurs on the port. This DDM functionality will be enabled on the other Ethernet Routing Switches in future software releases.

Review the fiber requirements of the network and select the appropriate GBIC/SFP/XFP based on those fiber specifications.

Fiber	62.5µ MMF		50µ MMF			9µ SMF	Wavelength
MHz*Km	160	200	400	500	2000	-	
10GBASE-SR	26m	33m	66m	82m	300m	-	850nm
10GBASE-LR/LW	-	-	-	-	-	10km	1310nm
10GBASE-ER/EW	-	-	-	-	-	40km	1550nm
10GBASE-ZR/ZW	-	-	-	-	-	80km	1550nm

Table 2.18: XFP Specifications

Transceiver	Speed	Fiber Type	Wavelength	Minimum Range	Maximum Suggested Range
1000SX	1 Gigabit	MMF – 62.5µ	850 nm	2-275 m	1.0 km
1000SX	1 Gigabit	MMF – 50µ	850 nm	2-550 m	1.0 km
1000LX *	1 Gigabit	MMF – 62.5µ	1310 nm	2-550 m	8.5 km
1000LX *	1 Gigabit	MMF – 50µ	1310 nm	2-550 m	8.5 km
1000LX	1 Gigabit	SMF – 9µ	1310 nm	2m-10 km	32.0 km
1000XD	1 Gigabit	SMF – 9µ	1550 nm	See Notes	50.0 km
1000ZX	1 Gigabit	SMF – 9µ	1550 nm	See Notes	70.0 km
CWDM	1 Gigabit	SMF – 9µ	1470-1610 nm	See Notes	See Note
* LX over MMF may require mode conditioning patch cables (single mode/multimode hybrid)					
Notes: 1000XD GBIC – if range is less than 25 km, use a 5 dB in-line attenuator 1000ZX GBIC – if range is less than 25 km, use a 10 dB in-line attenuator, and if range is less than 50 km, use a 5 dB in-line attenuator CWDM ranges vary with the SFP or GBIC that is used. Two versions of SFP are available (40 km and 70 km) and one version of GBIC (120 km). Please refer to product documentation for optical specifications and possible maximum ranges for CWDM.					

Table 2.19: GBIC / SFP Specifications

It is imperative to preserve the integrity of the uplinks from the edge closet to the core/distribution layer. In the case of a single fiber fault – either the transmit or the receive – the link must be automatically disabled at both ends. If this does not happen, data could be passed to a port that is not operating properly, and that data would be lost. This issue can cause severe performance degradation and eventually render the network inoperable. To protect the network, it is important to properly enable some form of single fiber fault detection on all uplink ports. VLACP can also be enabled as another single fiber fault detection mechanism.

All uplink ports should be configured with a method to detect a single fiber fault and disable the affected ports. There are two options for enabling this feature, depending on the switching platform being utilized at the edge:

- Switches supporting Autonegotiation on the uplink ports should have that feature enabled on both ends of the uplink.
    - On the Gigabit ports, Remote Fault Identification (RFI) is enabled by using Autonegotiation. RFI removes the link from a port in the event of a single fiber fault on the link connected to that port.
    - On 100 Mbps ports, Far End Fault Identification (FEFI) is enabled by using Autonegotiation. FEFI removes the link from a port in the event of a single fiber fault on the link connected to that port.
    - 10-Gigabit does not support autonegotiation; however, RFI is enabled as part of the physical layer implementation automatically.

In certain scenarios where the link may span across a providers LAN extension service, detecting a link failure in the LAN extension core will not work using Autonegotiation and RFI. This will not work end to end between a pair of Avaya switches because RFI only operates between direct connected switches. Hence, if there is a failure in the LAN extension core, the link on both Avaya switches will still be running. To solve this problem, enable VLACP on the Avaya switches. The VLACP protocol is forwarded between the Avaya switches. If the switch does not receive any VLACP updates, a link will be declared out-of-service and traffic will be forwarded through another working link. A detailed discussion of VLACP is covered in an upcoming section.

2.2.5 Autonegotiation

The autonegotiation standard for Ethernet allows end stations to connect at their most optimal speed and duplex – anything from 10 Mbps half-duplex up to 1 Gbps full duplex. This feature allows different end stations with different connectivity capabilities to connect to a single network without the intervention of the network administrator.

Autonegotiation must be enabled on fiber ports to support Remote Fault Identification (RFI) for Gigabit and Far End Fault Identification (FEFI) for 100FX connections. These features shut down a port in the case of a single fiber fault at the remote end of the link. A more detailed description is covered in the section on fiber fault detection.

10-Gigabit Ethernet does not have the concept of autonegotiation. By default, 10-Gigabit links are full duplex. The RFI component of autonegotiation is built into the 10-Gigabit physical interface and therefore is automatically enabled.

It is critical to verify that both ends of the link are capable of supporting autonegotiation. If one end is not able to support it, then autonegotiation must be disabled on both ends of the link. Having autonegotiation enabled on one end and disabled on the other end is a common configuration error that can cause severe performance degradation of that connection. Excessive FCS (Frame Check Sequence) errors on a port are a common indicator of a speed/duplex mismatch between the two devices.

In certain situations, it is useful to autonegotiate to a specific speed and duplex value by controlling which capabilities are being advertised from the switch. Avaya introduced the Custom Auto-Negotiation

Advertisement (CANA) feature to accommodate this need. This feature allows the administrator to control which advertisements are made by the switch. For example, if the switch only advertises a 100 Mbps full duplex capability on a specific link, then the link is only activated if the neighboring device is also capable of autonegotiating a 100 Mbps full duplex capability. This prevents mismatched speed/duplex modes if customers disable autonegotiation on the neighboring device.

- Enable autonegotiation on all switch-to-switch ports to ensure single fiber fault detection.
- Enable autonegotiation on all end station ports, ensuring that those stations are capable of supporting autonegotiation.
- Where required, use Customized Auto-Negotiation Advertisements (CANA) to control end station connectivity speed and duplex.
- Disable autonegotiation on problematic devices. Because the autonegotiation standard is rather broad, there are some devices that will not connect properly when autonegotiation is enabled on both ends.
- When using autonegotiation, always have the most recent Network Interface Card (NIC) driver from the manufacturer.

2.2.6 Link Aggregation

To increase both resiliency and bandwidth from the edge-wiring closet, Avaya recommends implementing link aggregation of the uplinks. Avaya supports multiple options for link aggregation on the Ethernet switches, MultiLink Trunking (MLT), Switch Clustering using Split MultiLink Trunking (SMLT), and 802.3ad.

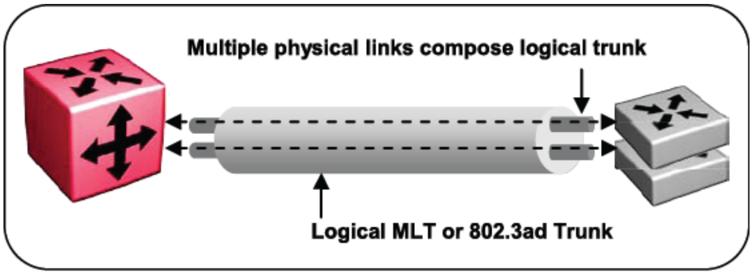
Multiple physical links compose logical trunk
Logical MLT or 802.3ad Trunk

Figure 2.33: Link Aggregation

MultiLink Trunking (MLT) provides the ability to group multiple physical links into a single logical link (see Table 2.20 for the specific number of links and groups supported per switch or stack). MLT automatically increases bandwidth from the wiring closet by utilizing all the physical links in a logical group. A failure of any physical links results in automatic sub-second failover of the traffic to the remaining links within the MLT group. After the failed link has been repaired, recovery of that link back into the MLT group is also accomplished in sub-second time.

Distributed MultiLink Trunking (DMLT) adds the ability to terminate the physical links of the trunk group on different switches within a stack or on different modules within a chassis. This feature significantly increases the resiliency of the uplinks out of the wiring closet. A failure of the switch or module on one of these physical links will not cut off communication from the wiring closet to the core/distribution layer.

The IEEE standard for link aggregation is 802.3ad. This standard uses the Link Aggregation Control Protocol (LACP) to aggregate multiple physical links into a single logical link aggregation group (LAG) from the Ethernet switch. Adherence to the IEEE standard will help to ensure interoperability of link aggregation between different vendor equipment (switches, NICs, etc.).

Although IEEE 802.3ad-based link aggregation and MLT provide similar services, MLT is statically defined, whereas IEEE 802.3ad-based link aggregation is dynamic and provides added functionality through the Link Aggregation Control Protocol (LACP). LACP dynamically detects whether links can be aggregated into a link aggregation group and does so when links become available. IEEE 802.3ad was designed for point-to-point link aggregation only. The Avaya Ethernet switches support the formation of 802.3ad link aggregation groups across different switches in a stack or different I/O modules in a chassis to provide additional resiliency in the event of a switch failure in a stack or an I/O module failure in a chassis.

The added functionality of end-to-end checking is a resiliency feature that should be used within a Converged Campus design. Because MLT is statically defined, there is no mechanism for checking between directly connected switches. In the rare event that a switch may become inoperable but the link status remains up, data could be inadvertently sent to that switch causing it to be lost. A checking mechanism automatically removes that port from the aggregation group, ensuring data is not sent to an inoperable switch, and therefore ensuring no data loss.

There are limitations to LACP that must be considered:

- LACP scaling can be somewhat limited, please see Table 2.20 below for Ethernet Switch support of LACP and scaling numbers
- LACP was designed to operate between two directly connected switches and will not work in an end-to-end fashion if there are any intermediary devices (Optical ring, Service Provider Network, etc.).
- Failover times of LACP are higher than that of MLT – default value is a 30 second poll and 3 polls must be missed to disable a port. Thus, it will take LACP 90 seconds to failover by default. Fast timers are available that reduce the poll time to 1 second intervals, thus reducing failover time to 3 seconds.

For the purposes of the Super Large Campus Solution, a MultiLink Trunk Group (MLT) will be created on the Edge switch/stack connecting to the VSP 9000 Switch Cluster core. Each closet containing edge switches/stacks will have an MLT created to the core of the network.

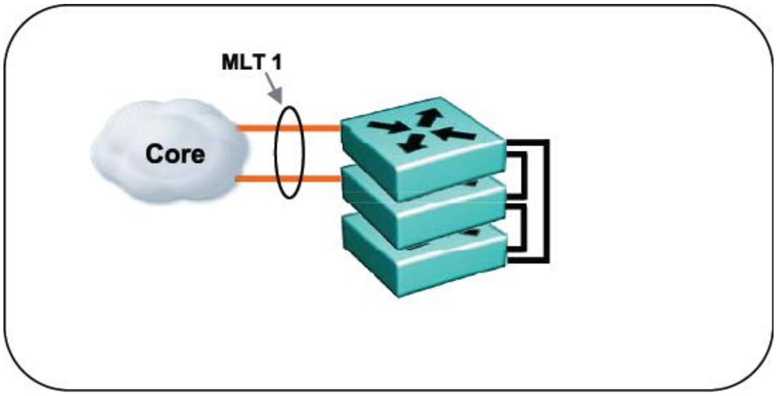

Figure 2.34: Edge Switch Link Aggregation

- Create an MLT group with id 1 and name it MLT-to-VSP 9000. The MLT id on the edge stack can be the same across the network. This id is only locally significant to that stack and is not propagated outside the local stack.
- Spanning Tree Protocol must be disabled for the MLT when using it with SMLT/SLT in the core of the network.

- The uplink ports will be assigned as members of the MLT group. The uplink ports are 802.1Q tagged and members of the Data, Voice, and Management VLANs, along with any other VLANs used at the Edge.
- When in a stack configuration at the edge, use DMLT (Distributed MLT). This adds another layer of resiliency to the design by spreading the MLT across different units in the stack.
- Autonegotiation should be enabled (on by default) on the uplink ports to enable Remote Fault Indication (prevention of single fiber faults)
- All links in an MLT group must be of the same speed and duplex – Avaya does not support the mixing of different speed links within the same MLT group.
- For specific configuration rules about link aggregation, see the product documentation.
- Avaya recommends using MLT whenever possible, however, if using LACP review the following:
  - LACP supports a maximum of eight active links, all other links (nine and above) are put into standby
  - Active/Standby is defined by ActorPortPriority (higher actor priority = lower port priority)
  - If actor priority is the same, lower MAC = higher priority
  - Use the same ID number for the ActorAdmin key and MLT ID
- In an RSMLT square/full mesh core topology, Avaya recommends to set the MLT algorithm on Avaya edge stackable switches to advance for IP based traffic. Tests have shown that this will increase performance over the default MLT setting of basic.

The following table depicts the support and scaling of MLT and 802.3ad on Avaya Ethernet switches.

Switch Model	Links per Group	Groups per Switch or Stack	802.3ad Support	LACP-MLT Scaling	LACP-SMLT Scaling
VSP 9000	16	512	Rls 3.0	511	511
ERS 8300	4	31*	Rls 4.1	31	31
ERS 5000	8	32	Rls 4.1	32	Future
ERS 4500	4	8	Rls 5.0	8 x 32	N/A
ERS 2500	4	6	Rls 4.0	6	N/A
* Up to seven Fast Ethernet groups and/or 31 Gigabit groups					

Table 2.20: LACP / VLACP Support and Scaling

2.2.7 VLANs

The ERS 5000 series supports up to 1024 VLANs while the ERS 4500 and ERS 2500 series support up to 256 VLANs. The following details the basic VLAN configuration parameters to consider and configure:

- Assign general use VLANs by geographic location if possible (by closet or floor). This practice limits the need to bridge VLANs throughout the network, thus reducing administrative overhead. It also aids in troubleshooting any network or application issues that may arise. An exception to this practice occurs when creating a VLAN for a specific protocol or application that may need to be bridged throughout the network.
- VLAN config control is strict – Globally enabled by default. In this mode, you must first remove port members from the default VLAN 1 prior to adding these port members to a new VLAN.
- Data VLAN and Voice VLAN to separate traffic. These VLANs can be modified to suit the specific environment needs and additional VLANs may be required for specific network services or user connectivity.
- Port Membership – All switch ports will be removed from the default VLAN (1) and added as members to the Data VLAN and Voice VLAN – the PVID for the ports must be the Data VLAN. See Section 2.2.13 Quality of Service for a discussion about dynamically creating the voice VLAN using Auto Detect Auto Config (ADAC) on the ERS stackables.
- Always enable 802.1Q VLAN tagging on uplink ports and have all VLANs (except VLAN 1) added to them. Even if only one VLAN is used at the edge, this enables the addition of more VLANs in the future without disrupting existing traffic. Enabling 802.1Q VLAN tagging adds a Q tag to every packet on the uplink in order to maintain the VLAN separation across the link.
- Avoid using the default VLAN whenever possible. This helps to minimize the possibility of accidentally creating Layer 2 loops in the network. Because the same default VLAN exists on every switch, it is very easy to incorrectly connect all these VLANs together and create unexpected traffic flows in the network. Please note that you cannot delete the default VLAN, therefore, it is recommended to remove valid port members from the default VLAN.
- Management – Create a separate management VLAN.

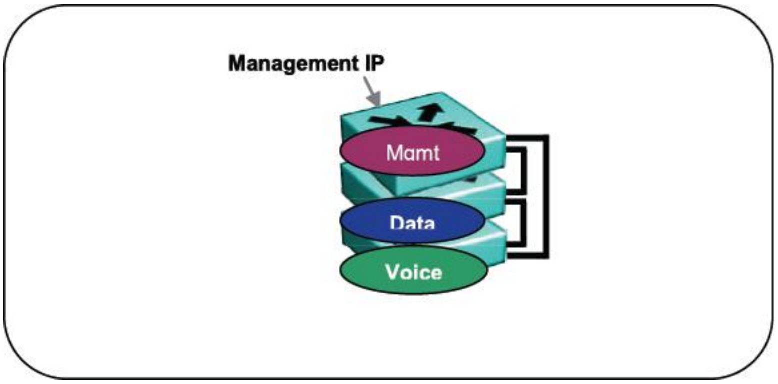

Figure 2.35: Edge Switch VLANs

2.2.8 Filter Untagged Frames

To provide protection against a factory defaulted or misconfigured device being connected to the uplink ports on the Edge switch/stack, which can result in a loop, the Filter Untagged Frames feature should be enabled on the MLT ports.

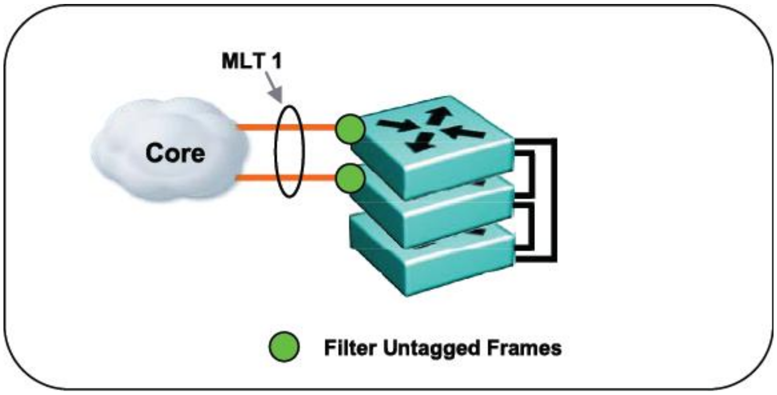

Figure 2.36: Edge Switch Filter Untagged Frames

2.2.9 Spanning Tree Protocol

Spanning Tree protocols detect and eliminate logical loops in a bridged or switched network. If multiple paths exist, the spanning tree algorithm configures the network so that a bridge or device uses only the most efficient path. If that path fails, the protocol automatically reconfigures the network and makes another path active, thereby sustaining network operations.

Virtual Services Platform 9000 supports RSTP and MSTP but can downgrade a port automatically if it receives an STP Bridge Protocol Data Unit (BPDU) from a switch that runs STP.

ⓘ Note: MSTP is the default spanning tree protocol on the VSP 9000.

Spanning Tree Groups (STG) represent logical topologies. A topology is created based on bridge configuration values such as root bridge priority. In the case of multiple STGs, you can map a VLAN to the most appropriate logical topology in the physical network. The default STG for Virtual Services Platform 9000 is STG 1. In RSTP mode all VLANs run in the default STG. In MSTP mode, you can create additional STGs by using the VLAN create command. The Virtual Services Platform 9000 supports up to 64 STGs.

When you map VLANs to STGs, be aware that all links on the bridge belong to all STGs. Because each Spanning Tree group can differ in its decision to make a link forwarding or blocking, you must ensure that the ports you add to a VLAN are in the expected state. Untagged ports can only belong to one VLAN and therefore to only one STG. Tagged ports can belong to multiple VLANs and therefore to multiple STGs. If a tagged port belongs to more than one STG, the spanning tree BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.

The Rapid Spanning Tree Protocol (RSTP or IEEE 802.1w) reduces the recovery time after a network breakdown. It also maintains backward compatibility with IEEE 802.1d (the spanning tree implementation prior to RSTP). In certain configurations, the recovery time of RSTP can be reduced to less than 1 second. RSTP also reduces the amount of flooding in the network by enhancing the way Topology Change Notification (TCN) packets are generated.

With Multiple Spanning Tree Protocol (MSTP or IEEE 802.1s), you can configure multiple instances of RSTP on the same device. Each RSTP instance can include one or more VLANs. By using RSTP and MSTP, the Virtual Services Platform 9000 achieves the following:

- reduce convergence time after a topology change (from 30 seconds to less than 1 or 2 seconds)
- eliminate unnecessary flushing of the MAC database and the flooding of traffic to the network
- create backward compatibility with classic 802.1d switches
- create support for 64 instances of spanning tree in MSTP mode

Avaya recommends using the Spanning Tree protocol on all end station connections in order to safeguard the network from hubs or other devices that could be inserted into the network at the end station. A modification to the normal learning of Spanning Tree is available in all Avaya edge switches. This feature is known as Fast Start or Fast Learning, and is the recommended setting for all end station ports.

The BPDU filtering feature also adds a level of protection against inadvertent loops in the network. This feature was originally developed to prevent an unwanted root selection process when a new device was added to a Spanning Tree network and/or to prevent unknown devices from influencing an existing spanning tree topology. A more detailed discussion on BPDU filtering follows in the next section.

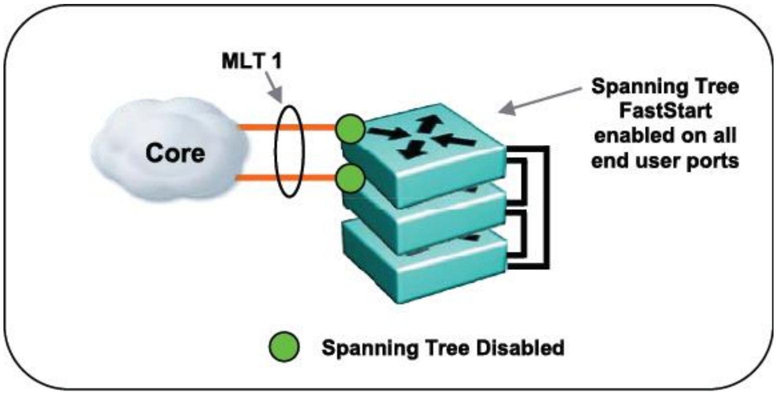

Figure 2.37: Edge Switch Spanning Tree

- Enable Spanning Tree Fast Start/Fast Learning on all end station ports.
- Enable BPDU filtering on all end station ports connected non network infrastructure devices such as PC’s, printers and IP phones.
- Never enable Fast Start/Learning on any uplink ports; this could cause unexpected behaviors on the entire network.
- When using Spanning Tree, pay attention to the root bridge. Ensure the root bridge is one of the core switches by configuring the Spanning Tree priority.
- When using SMLT to connect the edge to the distribution/core, always disable Spanning Tree on the uplink ports/MLT of the edge switch.

Avaya recommends using Split MultiLink Trunking (SMLT) to interconnect closets to the core of the network, thus eliminating the need for the Spanning Tree protocol on uplinks. When using SMLT between the edge switch and the core or distribution switch, two or more redundant paths to two separate core/distribution switches are utilized in an active-active fashion without the need for Spanning Tree to prevent loops. Traffic is distributed over all available paths using either MLT, 802.3ad, or any other form of link aggregation. If one or more of the paths fail, including link and/or switch failures, SMLT provides sub-second failover to the remaining path(s).

2.2.10 BPDU Filtering

BPDU filtering is normally used as a protection mechanism within a Spanning Tree network. This feature blocks an unwanted root selection process when a device is added to the network and also blocks BPDU packets from ingressing the port. This feature is also a valuable protection mechanism in a Switch Cluster network as it will block loops created by connecting Edge closets directly together, therefore it is recommended to enable BPDU Filtering on all end user ports.

- ➢ Set Timeout to 0 – requires manual intervention to re-enable the port

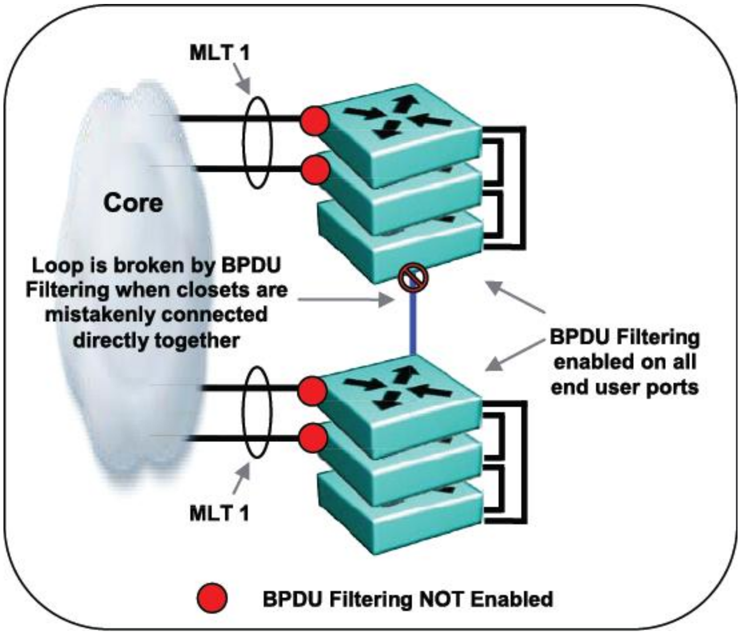
MLT 1
Core
Loop is broken by BPDU Filtering when closets are mistakenly connected directly together
BPDU Filtering enabled on all end user ports
MLT 1
● BPDU Filtering NOT Enabled

Figure 2.38: Edge Switch BPDU Filtering

2.2.11 VLACP

This Avaya feature provides an end-to-end failure detection mechanism which will help to prevent potential problems caused by misconfigurations, a switch being defaulted, or links connected incorrectly. In the Super Large Campus design, VLACP will be used between the VSP 9000 Switch Cluster and all uplink ports to the ERS 2500/4500/5000 switches at the edge.

- ➢ Globally configure VLACP to use the reserved multicast MAC of 01-80-C2-00-00-0F
- ➢ For the SLT/SMLT links, use the short timeout
  - ▪ Fast periodic timer of 500 msecs * timeout scale of 5
  - ▪ Make sure these values match on both ends of the links

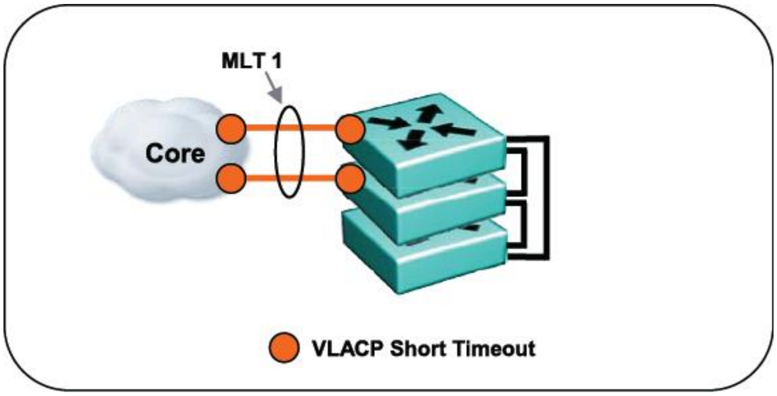

Figure 2.39: Edge Closet VLACP

2.2.12 Rate Limiting

Port level rate limiting limits packets with broadcast and/or multicast addresses to control the amount of traffic on a port. Rate limiting is used to protect the network from non-CPU bound traffic. This functionality is configured on a per port basis. For each port, the network administrator can configure a rate limit for broadcast traffic and a rate limit for multicast traffic. These rates are the maximum allowed amount of that traffic type on that specific port. When traffic exceeds the configured threshold, it is dropped. The design recommendations below detail the way broadcast/multicast traffic is calculated (%, pps, kbps) and the values that should be used. Be aware that these are rule-of-thumb values for the "typical" enterprise network.

It is extremely important to understand the network and application environment before configuring the rate limiting feature. In certain environments, there will naturally be a higher rate of a traffic type due to the applications being used. For example, in a network that utilizes multimedia communications such as streaming video and video on demand, there will likely be a higher rate of multicast traffic and rate limiting this traffic may adversely affect the applications being used.

If rate limiting is implemented, it should be done at the edge of the network closest to the user as possible. This will have the greatest effect on overall traffic as limiting will occur before hitting the core of the network.

The following details the implementation of rate limiting on each of the ERS platforms and the recommended values to be used.

ERS 2500 / 4500 / 5000

- 1 – 10% of port speed
- Recommendation → 10%

ERS 8300

- 1-100% of port speed
- Recommendation → 10%

VSP 9000

- Broadcast / multicast bandwidth limiting
- Allowed rate is in kbps
- Recommendation → 3 times normal kbps

2.2.13 Quality of Service

To provide appropriate QoS treatment for Voice traffic, DiffServ will be enabled on the Edge Switches/Stacks. The QoS capabilities of the Edge switching platforms are detailed below. Although Voice traffic is normally the driver for a QoS implementation, any real time delay intolerance or business critical traffic can take advantage of QoS over the network.

From a generic QoS perspective, the Edge switches support the following:

The ERS 2500 presently can be configured to honor DSCP markings within the packets entering the switch and placing them in the appropriate egress queue.

- Configure the ERS 2500 to honor DSCP markings on all ports to ensure proper QoS.

The ERS 4500 and ERS 5000 series have advanced QoS features and can be configured to mark and/or honor DSCP markings within the packets entering the switch and place them in the appropriate egress queue.

- Configure the ERS 4500 or ERS 5000 to prioritize any traffic type that requires QoS throughout the network. This can be done by creating filters to identify the traffic (tcp port, udp port, src/dst IP, etc.) and then apply the appropriate DSCP marking to the packets.

When deploying IP Telephony, Avaya provides different options for the configuration of the Ethernet switches to ease the deployment and later moves, adds, and changes that occur within the normal Enterprise network. Each of these options will be described briefly here and references to Solution Guides and Configuration Guides are provided for more in-depth coverage.

- The ADAC (Auto Detect Auto Config) feature will automatically discover the IP phone once plugged into the Ethernet edge switch. Once discovered, the Ethernet edge switch will automatically provision the appropriate Voice VLAN and QoS to the port as well as on the uplink to the core.

  ADAC will use 802.1AB LLDP or the phone's MAC address as the discovery mechanism. This fact allows ADAC to be IP phone agnostic and will work with any 3rd party device.

  For a detailed discussion of ADAC and VoIP deployment, please refer to the ***Avaya IP Phone Set InterWorking with Avaya ES and ERS Switches TCG (NN48500-517)***.

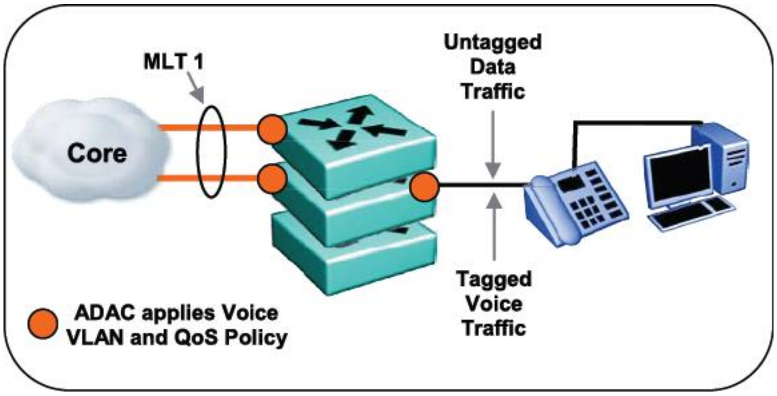

Figure 2.40: ADAC in Tagged Frames Mode

- The Avaya Automatic QoS feature will automatically discover the IP phone once plugged into the Ethernet edge switch. Once discovered, the Ethernet edge switch will automatically honor the QoS marking (DSCP) of the Avaya IP phone. Please note that Avaya Automatic QoS is only supported on a limited number of heritage Avaya IP Phones and call servers.

  A big advantage of Avaya Automatic QoS is its ability to recognize the 2050 Softphone and provide QoS treatment at the edge for these devices.

  Please note that if Voice VLANs are required they will need to be manually configured on the Ethernet edge switches, both for the end stations and the uplinks.

  For a detailed discussion of ADAC and VoIP deployment, please refer to the ***Avaya Automatic QoS Technical Configuration Guide (NN48500-576)***.

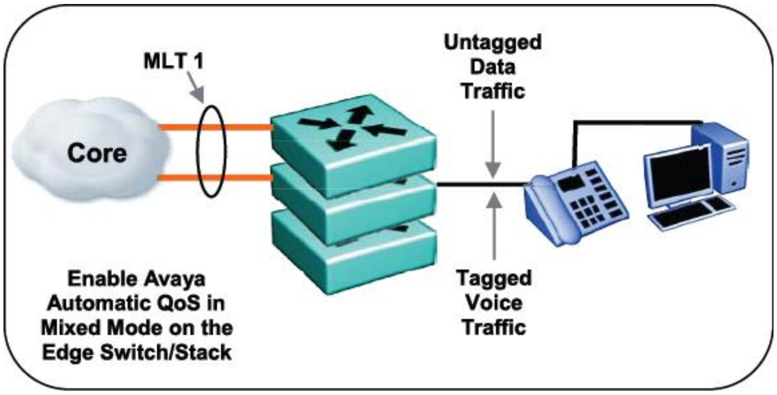
MLT 1
Untagged Data Traffic
Core
Enable Avaya Automatic QoS in Mixed Mode on the Edge Switch/Stack
Tagged Voice Traffic

Figure 2.41: Avaya Automatic QoS

2.2.14 Security

The Edge switches support security features to help protect the network from denial of service (DoS) attacks. DHCP Snooping, Dynamic ARP Inspection, and IP Source Guard should be enabled on all edge switches to help ensure the protection and integrity of the Campus network. See Section 2.5 for a more detailed discussion about security features.

- DHCP Snooping prevents DHCP spoofing by creating a DHCP binding table to help ensure that no rogue DHCP servers can be inserted into the network.
- Dynamic ARP Inspection examines ARP packets to prevent man-in-the-middle attacks.
- IP Source Guard filters clients with invalid IP addresses.
- Untrusted ports are normally the end user access ports on the switch.
- Trusted ports are normally the uplinks from the edge switch to the core.
- Dynamic ARP Inspection and IP Source Guard use the binding table created by DHCP Snooping, therefore, in order to use these features, DHCP Snooping must be enabled.
- IP Source Guard should not be enabled on uplink ports from the Edge to the Core, only enable on Edge access ports (untrusted ports) where DHCP Snooping and Dynamic ARP Inspection are enabled.

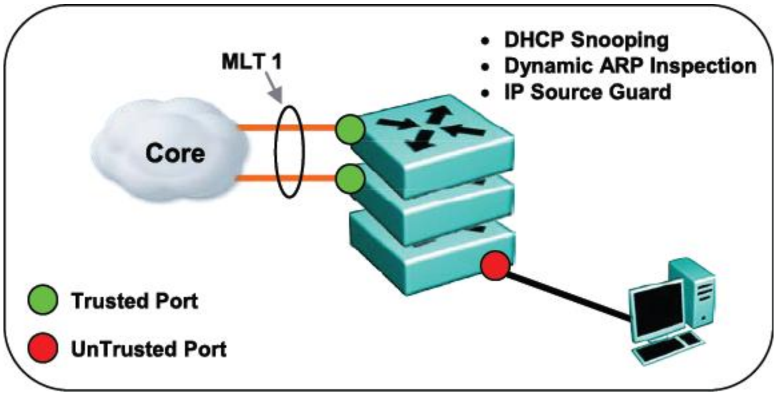

Figure 2.42: Edge Switch Security

2.2.15 Multicast

The multicast protocol distributes traffic to all subscribers of a multicast group. The Edge Switch/Stack is functioning at Layer 2 in this design and therefore IGMP features must be properly configured for the most efficient use of the bandwidth available.

- IGMP Snooping

  The Avaya Ethernet switches can sense IGMP (Internet Group Multicast Protocol) host membership requests for each specific multicast group. Only host ports requesting a multicast stream receive that stream; the switch automatically prunes the other ports and does not send multicast traffic to hosts that did not request it, thus making efficient use of the available bandwidth to each of the hosts on the switch.

- IGMP Proxy

  The Avaya Ethernet switch provides a single proxy report upstream for all members within the same multicast group on the same switch/stack. By consolidating all the IGMP host membership requests into a single request, the switch does not flood the network needlessly with multiple copies of the same request. IGMP Snooping must be enabled for this feature to work.

- IGMP Static Router Port

  This feature allows unknown multicast traffic to be forwarded only to the statically defined multicast router port. The traffic will not be flooded to all ports and will not be sent to dynamically-learned multicast router ports. Static router ports serve two purposes: to get (1) multicast traffic and (2) IGMP reports to multicast routers that may not be discoverable through passive detection; for example, when there are two queriers and one is elected and the other becomes silent (per the IGMP standard).

If IGMP Snooping/Proxy is not enabled, multicast traffic is flooded to all ports on the switch that are in the same VLAN.

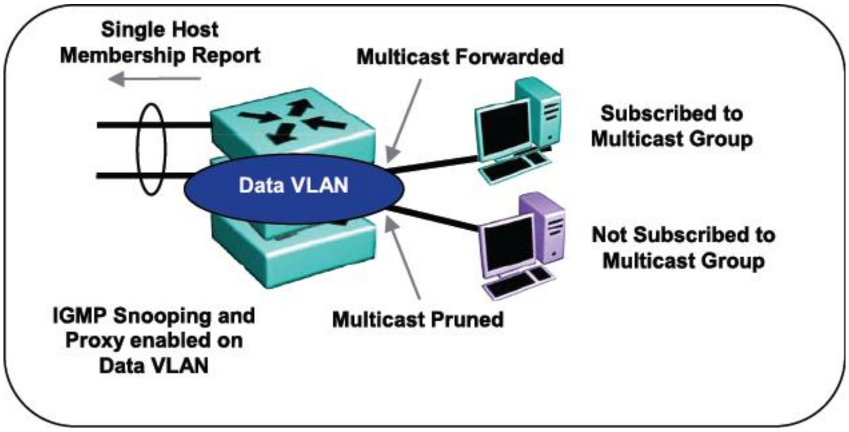
Single Host Membership Report
Multicast Forwarded
Subscribed to Multicast Group
Data VLAN
Not Subscribed to Multicast Group
IGMP Snooping and Proxy enabled on Data VLAN
Multicast Pruned

Figure 2.43: Edge Switch Multicast

- Enable IGMP Snooping on VLANs that have Multicast applications running on them – this prunes multicast traffic from end user ports that are not subscribed to the multicast group.
- Enable IGMP Proxy on the VLANs that have Multicast applications running on them – this will consolidate IGMP reports into a single report to the upstream network and is thus a much more efficient use of the bandwidth.
- IGMP Proxy must have IGMP Snooping enabled before it takes effect.
- IGMP Snooping and Proxy require an IGMP Querier to be present on the network. This function is performed by the Layer 3 Router with a Multicast routing protocol enabled or with the Layer 2 multicast querier function available on the ERS 8300.
- If there are no Multicast applications being used on the network, leave IGMP snooping and IGMP proxy disabled (default value) on all the VLANs.
- Flooding unknown multicast traffic is a behavior that is configurable (on or off) on some models of the Ethernet edge switches. The `vlan igmp unknown-mcast-no-flood` command provides this functionality. Turning off flooding ensures that static router ports become the destination of unknown multicast traffic.

2.3 Network Access Control

Avaya's Identity Engines is the framework for role-based network access control. Within this framework there are several options to best accommodate the needs of the Enterprise customer, from simple MAC authentication to full 802.1X authentication and posture assessment (end station compliancy to corporate security policies) of the end user's workstation. With all the methods available, the end result is to ensure users are allowed on the network and permitted access to resources based on identity and credentials.

This section will describe the backend infrastructure required (Identity Engines) along with the options available for end user authentication.

2.3.1 Identity Engines

The Avaya Identity Engines portfolio integrates with any current network infrastructures to provide the central policy decision needed to enforce role-based Network Access Control (NAC). This is accomplished by combining the best elements of a next-generation RADIUS/AAA server, the deep directory integration found in application identity offerings, and one of the industry's most advanced standards-based policy engines. All this is done out-of-band for maximum scalability and cost effectiveness.

The centralized policy engine sits in the data center to provide centralized authentication and authorization for wired, wireless, and VPN network devices. It is closely aligned with Avaya and third-party Ethernet switching, WLAN and VPN products as it provides centralized integrated security services for these network devices.

Coupled with the centralized policy engine is a suite of complementary products that enable 802.1X rollouts for wired and wireless networks, while unifying those policies with existing VPN rules to achieve audit and compliance goals. These products offer a holistic network identity management solution which involves all aspects of managing how users access networks. Benefits include admission control, temporary user provisioning, policy decisions and directory integration.

2.3.1.1 Identity Engines Ignition Server

A state-of-the-art network identity management solution with a powerful policy engine to centralize, streamline and secure access across the network. The Identity Engines Ignition Server offers a new level of accuracy, with identity- and policy-based control over who accesses the network, where, when, how, and with what type of device. Easy to deploy and use, it is a powerful, scalable foundation for network access control, guest access, secure wireless, compliance, and more.

2.3.1.2 Identity Engines Ignition Posture

Identity Engines Ignition Posture provides endpoint health and posture checking that works in the real world. Most posture checking products today are inflexible, add-on layers that are expensive to support and frustrating for network users. In contrast, this product provides policy flexibility and integration with the Identity Engines Ignition Server to ensure that it is easier to support and less frustrating for users.

2.3.1.3 Identity Engines Ignition Guest Manager

Because guests and visitors often have legitimate reasons to access networks, Identity Engines Ignition Guest Manager makes it easy and safe for organizations to let front-desk staff create guest user accounts. Simple delegation rules ensure front-desk personnel can give guests access to only specified network resources, and each guest account expires automatically after a designated period.

2.3.1.4 Identity Engines Ignition Analytics

Identity Engines Ignition Analytics is a powerful reporting application that allows organizations to perform in-depth analysis of network activity including ingress and usage. With over 25 preconfigured audit, compliance and usage reports, organizations can easily produce multiple custom reports to fulfill its specific reporting requirements.

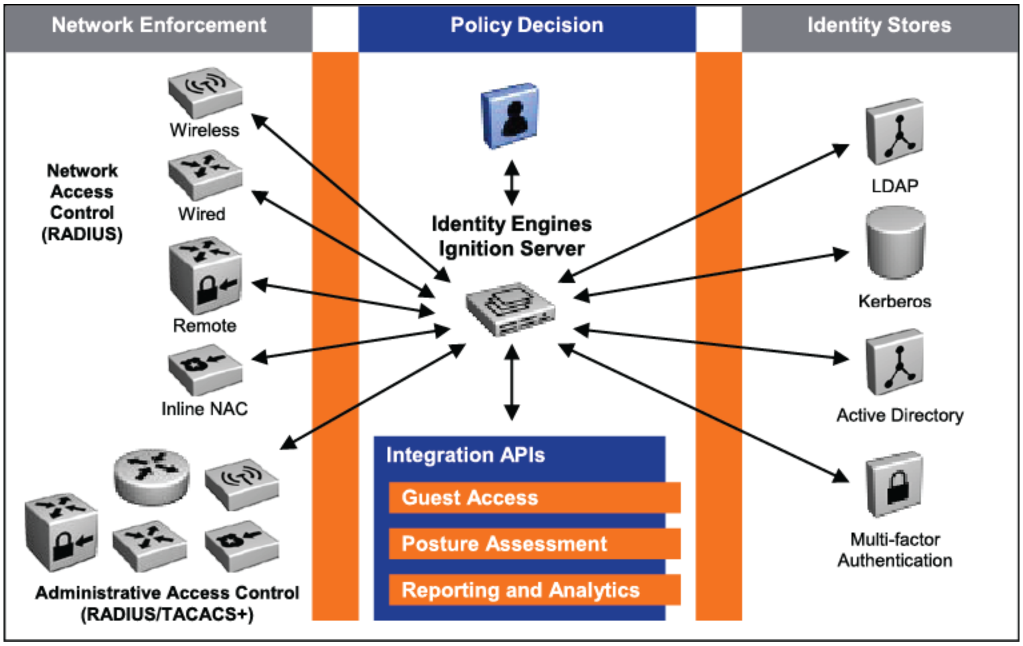

Figure 2.44: Identity Engines Portfolio Architecture

2.3.2 MAC Based Authentication

The Media Access Control (MAC) address-based security feature is based on Avaya BaySecure local area network (LAN) Access for Ethernet, a real-time security system that safeguards Ethernet networks from unauthorized surveillance and intrusion. You can use the MAC address-based security feature to configure network access control, based on the source MAC addresses of authorized stations.

- Local Authentication

  The MAC Address security feature allows the administrator to specify a list of MAC addresses that are authorized to access the switch. You can also specify the ports that each MAC address is allowed to access. The options for allowed port access include: NONE, ALL, and single or multiple ports that are specified in a list.

- Auto-learning or Static Authentication

  Auto-learning allows a switch to automatically add a MAC address to the MAC security table without user intervention. You can also limit the number of MAC addresses allowed per port.

  Static authentication allows the manual creation of static MAC entries on a per port basis.

- Centralized Authentication

  This feature functions the same way as local authentication; however, the list of allowed MAC addresses is stored in a RADIUS server. This is a much more manageable approach to MAC security. Dynamic VLAN assignment is supported for MAC authenticated clients.

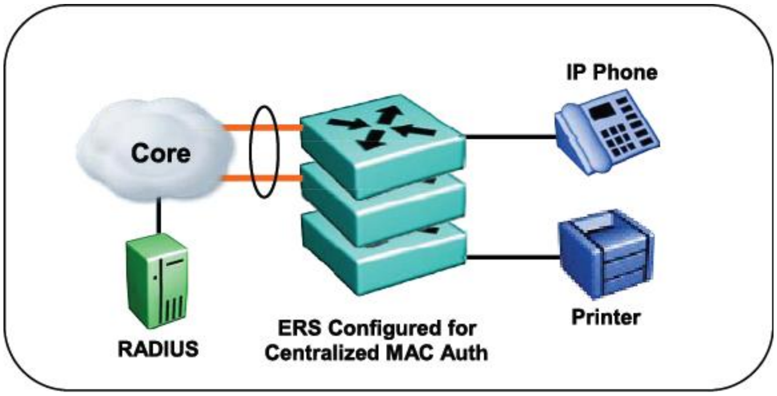

Figure 2.45: MAC Based Authentication

2.3.3 802.1X Extensible Authentication over LAN (EAPoL)

The Ethernet Routing Switches use an encapsulation mechanism to provide security, referred to as the Extensible Authentication Protocol over LAN (EAPOL). This concept uses the Extensible Authentication Protocol (EAP) as described in the IEEE 802.1X standard allowing the set up of network access control on internal LANs. EAP allows the exchange of authentication information between an end station or server connected to the switch and an authentication server, such as a RADIUS server. The EAPOL-based security feature operates in conjunction with a RADIUS-based server to extend the benefits of remote authentication to internal LAN clients.

A variation of 802.1X allows for the configuration of a guest VLAN, which allows non-authenticated users on the network, but also allows the administrator to control what the guest VLAN users can access. Other variations of 802.1X allow for multiple authenticated users on the same physical Ethernet switch port, or the ability to support both 802.1X and MAC authenticated users on the same port. This flexibility in using 802.1X makes it much easier to deploy scenarios with IP Telephony where the PC is connected to the IP phone and they share a single Ethernet switch port in the closet.

- Guest VLAN

  Allows non-EAP users connected on EAP-enabled ports access to a guest network. This feature allows network access to users through the guest VLAN. This VLAN is port-based, configured on a per switch/stack basis, and is enabled per port. Once a user completes EAP authentication, the user is moved from the Guest VLAN to the authenticated VLAN.

- Single Host Single Authentication (SHSA)

  For an EAP-enabled port with SHSA, at any time only one client (single MAC address) is authenticated on a port, which is assigned to only one port-based VLAN. Only a particular client who completes EAP negotiations on the port is allowed access to that port for traffic.

  - Dynamic VLAN assignment for the client is supported by RADIUS attributes passed back to the Ethernet switch.
  - Guest VLAN is supported in conjunction with SHSA.

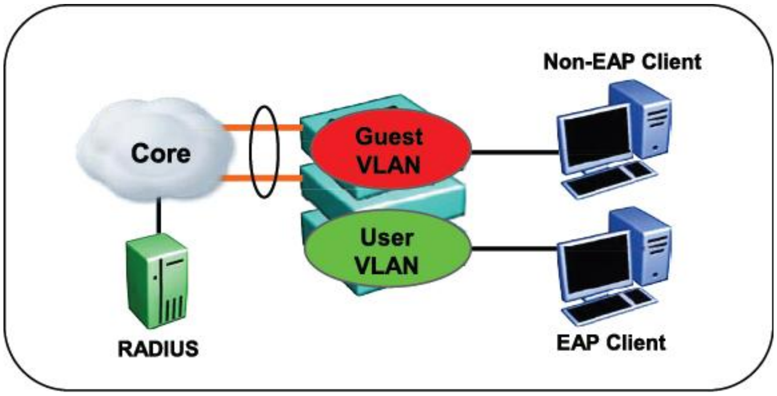

Figure 2.46: 802.1X SHSA

- Multiple Host Multiple Authentication (MHMA)

  For an EAP-enabled port with MHMA, a finite number of clients or devices with unique MAC addresses are allowed access to a port. Each client must complete EAP authentication to enable the port to allow traffic from the corresponding MAC address. Only traffic from the authorized hosts is allowed on that port.

  - Dynamic VLAN assignment for the client is supported by RADIUS attributes passed back to the Ethernet switch – note that the dynamic VLAN is supported for the first authenticated client only, subsequent client authentications will automatically be put into that initial dynamically assigned VLAN.
  - Guest VLAN is supported in conjunction with MHMA – after first successful EAP authentication, the Guest VLAN is no longer accessible on that port.

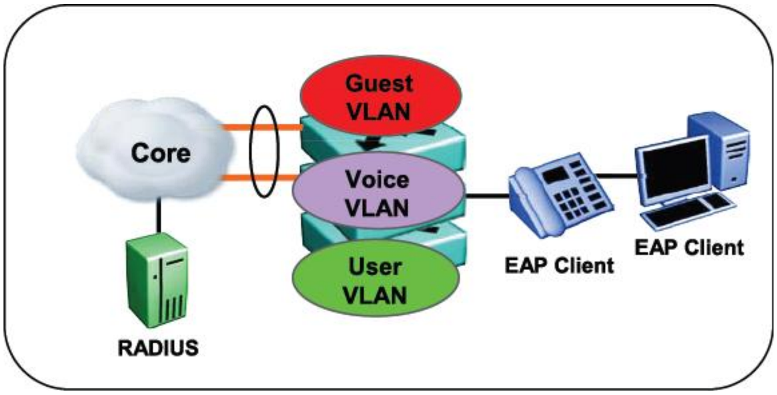

Figure 2.47: 802.1X MHMA

- Non-EAP MAC Authentication

  Allows non-EAP users connected on EAP-enabled ports access to the network. This feature uses MAC authentication for the non-EAP user. This authentication can be local MAC authentication on the Ethernet switch or centralized MAC authentication via RADIUS – depending on Ethernet switch feature support. Please note that user-based policies can be used with MAC Authentication. The non-EAP hosts are permitted on the port even if there are no EAP authenticated hosts on the port.

  - Guest VLAN and Non-EAP are mutually exclusive (ERS 4500 with Release 5.3 software removes this limitation)

- Non-EAP IP Phone Authentication

  The “non-eap-phone-enable” feature allows Avaya IP phones (heritage Nortel versions only) on EAP-enabled ports without the need for the phone MAC to be pre-configured in the MAC list (local or centralized). The switch will look for the phone signature enclosed in the DHCP Discover packet. If the proper signature is found, the switch will register the MAC address of the IP Phone as an authenticated MAC address and will let the phone traffic pass on the port.

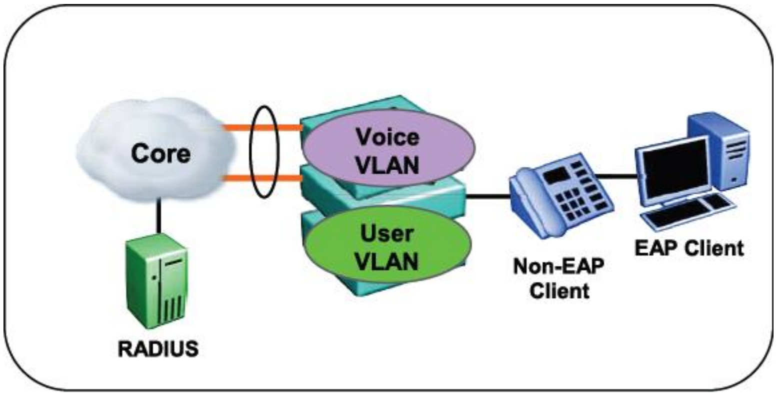

Figure 2.48: 802.1X Non-EAP Phone Authentication

- Multiple Host Single Authentication (MHSA)

  Multiple Host Single Authentication (MHSA) is a more restrictive implementation of support for non-EAPOL hosts on EAPOL-enabled ports. For an EAPOL-enabled port configured for MHSA, one EAPOL user must successfully authenticate before a finite number of non-EAPOL users or devices with unique MAC addresses can access the port without authentication. The MHSA feature is intended primarily to accommodate printers and other passive devices sharing a hub with EAPOL clients.

Authentication Features	ERS 2500 V 4.4	ERS 4500 V 5.5	ERS 5000 V 6.2
Single Host Single Authentication (SHSA)	Y	Y	Y
Multi Host Single Authentication (MHSA)	Y	Y	Y
Multi Host Multi Authentication (MHMA)	Y	Y	Y
Guest VLAN (GV)	Y	Y	Y

RADIUS Assigned VLANs for EAP MACs	Y	Y	Y
Port-mirroring on 802.1x (EAP) ports	N	Y	Y
EAP/NEAP Separation	N	Y	Y
User Based Policies (UBP)	N	N	Y
User Based Policies v2 (UBPv2)	N	N	Y
NEAP User Based Policies	N	N	Y
NonEap Support on EAP enabled ports (NEAP)	Y	Y	Y
IP-Phone Authentication based on DHCP Signature	Y	Y	Y
RADIUS Assigned VLANs for NonEap MACs	Y	Y	Y
NEAP & Guest VLAN on same port (VOIP VLANs)	Y	Y	Y
EAP/NEAP with Fail-Open VLAN (FOV)	Y	Y	Y
EAP/NEAP Last Assigned VLAN	Y	Y	Y
EAP/NEAP MultiVLAN Capabilities	N	Y	Y
EAP Accounting	Y	Y	Y
RFC 3576 RADIUS Request Server	Y	Y	Y
RADIUS Interim Accounting Updates	Y	Y	N
EAP/NEAP with VLAN names	Y	Y	Y
NEAP Re-authentication	Y	Y	N
RADIUS EAP NEAP different servers	Y	Y	N
NEAP RADIUS Accounting	Y	Y	N
RADIUS Management Accounting	N	Y	N
Radius Request use Management IP	Y	Y	Y
RADIUS Timeout Mode	N	Y	N
RADIUS Management Accounting Tacacs Support	N	Y	N

Table 2.21: Supported Authentication Features

2.4 Troubleshooting and Monitoring

Understanding what is happening during the normal course of operations and knowing what to look for during abnormal times can help to maintain connectivity or restore operations quickly. This section highlights a few critical and often used troubleshooting tools. For details on all the options available, refer to the Troubleshooting documentation for each Avaya product.

2.4.1 Packet Capture (PCAP)

The VSP 9000 and ERS 8600/8800 support a Packet Capture Tool (PCAP) tool that captures ingress and egress packets on selected I/O ports. With this feature, you can capture, save, and download one or more traffic flows through the VSP 9000. The captured packets can then be analyzed offline for troubleshooting purposes. This feature is based on the mirroring capabilities of the I/O ports. To use PCAP, you must have the Advanced Software License. All captured packets are stored in the Secondary CPU, used as the PCAP engine. The Master CPU maintains its protocol handling and is not affected by any capture activity.

2.4.2 Port Mirroring

The VSP 9000 and Ethernet Routing Switches offer a port mirroring feature that helps you monitor and analyze network traffic. Port mirroring supports both ingress (incoming traffic) and egress (outgoing traffic) port mirroring. When you enable port mirroring, ingress or egress packets are forwarded normally from the mirrored (source) port, and a copy of the packets is sent to the mirroring (destination) port. Port mirroring capabilities and scalability vary between platforms.

As a general rule, you can mirror different speed ports and different physical media

- 10/100 to Gig – no issues
- Gig to 10/100 – may not see all packets if exceeding 100 Mbps on the Gig link
- Copper to Fiber – no issues
- Fiber to Copper – no issues

Table 2.22 details the port mirroring capabilities of the ERS stackable portfolio.

Ethernet Switch	Port Mirroring	Mirrored Ports (Source)					Mirroring Ports (Dest)
		Total Number	Ingress	Egress	Both	MAC-based	
ERS 2500	1-to-1	1	Yes	Yes	Yes	No	1
ERS 4500	Many to 1	Many	Yes	Yes	Yes	Yes	1
ERS 5500	Many to 1	Many	Yes	Yes	Yes	Yes	1
ERS 5600	Many to Many	Many	Yes	Yes	Yes	Yes	1
Many = all ports on a switch or stack Many to Many applies to ERS 5600 hardware only and is implemented as four instances of Many to 1 mirroring. This functionality works only on ERS 5600 switches or pure stacks and will not work in a hybrid stack of ERS 5600s and ERS 5500s.							

Table 2.22: Stackables Port Mirroring Capabilities

Table 2.23 details the port mirroring capabilities of the VSP 9000 and the ERS modular portfolio.

Ethernet Switch	Port Mirroring	Mirrored Ports (Source)					Mirroring Ports (Dest)
		Total Number	Ingress	Egress	Both	MAC-based	
ERS 8300	Many to 1	Many	Yes	Yes	Yes	No	1
ERS 8600 Legacy Modules	Many to 1	Many	Yes	Yes	Yes	Yes	64
ERS 8600 R Modules (4.0/4.1)	Many to 1	1 per LANE	Yes	Yes	Yes	Yes	Many
ERS 8600 R Modules (5.0 and later)	1 to 1	1	1	1	Yes	Yes	1
	1 to Many	1	1	1	Yes	Yes	Many
	1 to MLT	1	1	Not Supported	Yes	Yes	MLT
	Many to 1	See Ingress/Egress	20 ports per LANE	1 port per LANE	Yes	Yes	1
	Many to Many	See Ingress/Egress	20 ports per LANE	1 port per LANE	Yes	Yes	Many
	Many to MLT	See Ingress/Egress	20 ports per LANE	Not Supported	Yes	Yes	MLT
ERS 8600 RS Modules (5.0 and later)	1 to 1	1	1	1	Yes	Yes	1
	1 to Many	1	1	1	Yes	Yes	Many
	1 to MLT	1	1	1	Yes	Yes	MLT
	Many to 1	See Ingress/Egress	20 ports per LANE	20 ports per LANE	Yes	Yes	1
	Many to Many	See Ingress/Egress	20 ports per LANE	20 ports per LANE	Yes	Yes	Many

	Many to MLT	See Ingress/Egress	20 ports per LANE	1	Yes	Yes	MLT
VSP 9000	1 to 1	1	1	1	Yes	Yes	1
	1 to Many	1	1	1	Yes	Yes	Many
	1 to MLT	1	1	1	Yes	Yes	MLT
	Many to 1	Many	No restrictions per LANE	No restrictions per LANE	Yes	Yes	1
	Many to Many	Many	No restrictions per LANE	No restrictions per LANE	Yes	Yes	Many
	Many to MLT	Many	No restrictions per LANE	No restrictions per LANE	Yes	Yes	MLT
1 to Many = mirroring of one port to all ports in a destination VLAN 1 to MLT = mirroring of one port to all ports in a destination MLT Many to Many = mirroring of many ports to all ports in a destination VLAN Many to MLT = mirroring of many ports to all ports in a destination MLT							

Table 2.23: Modular Port Mirroring Capabilities

Table 2.24 and Table 2.24 show the assignments of ports within an ERS 8800/8600 switch to the Octapids (for E/M modules) and Lanes (for R and RS modules). Note that ports belonging to the same Octapid group must be mirrored to the same destination port.

ERS 8800 Modules	Ports/Octapid	Port Assignments/Octapid (8 Octapids/Module)
8608 (GBE, GTE, SXE) 8608 (GBM, GTM)	1	1, 2, 3, 4, 5, 6, 7, 8
8616 (SXE, GTE)	2	1-2, 3-4, 5-6, 7-8, 9-10, 11-12, 13-14, 15-16
8624FXE	8	1-8, 9-16, 17-24
8632 (TXE, TXM)	8 per copper 1 per GBIC	1-8, 9-16, 17-24, 25-32, 33 (GBIC), 34 (GBIC)
8648 (TXE, TXM)	8	1-8, 9-16, 17-24, 25-32, 33-40, 41-48
8672 (ATME, ATMM)	4 with OC3 2 with DS3 1 with OC12	1-4, 5-8 with OC3 1-2, 3-4 with DS3 1,2 with OC12
8683POSM	2 with OC3 1 with OC12	1-2, 3-4, 5-6 with OC3 1, 2, 3 with OC12
8681XLR / 8681XLW	1 port uses all 8 Octapids	1

Table 2.24: ERS 8800 I/O Module Port to Octapid Mapping

ERS 8800 Modules	Ports/Lane	Port Assignments/Lane (3 Lanes/Module)
8630GBR	10	1-10, 11-20, 21-30
8648GTR	24	1-24, 25-48 (only uses 2 Lanes)
8683XLR / 8683XZR	1	1, 2, 3
8648GBRS	16	1-16, 17-32, 33-48
8648GTRS	24	1-24, 25-48 (only uses 2 Lanes)
8634XGRS	16, 16, 2	1-16, 17-32, 33-34
8612XLRS	4	1-4, 5-8, 9-12

Table 2.25: ERS 8800 I/O Module Port to Lane Mapping

2.4.3 Remote Logging

All of the ERS platforms support a remote logging feature. This provides an enhanced level of logging by replicating system messages on a syslog server. System log messages from several switches can be collected at a central location, alleviating the network manager from querying each switch individually to interrogate the log files. It also ensures that information is not lost when a switch becomes inoperable. The level of logging and the details provided differ between ERS platforms – refer to the System Monitoring or Troubleshooting documentation for each of the products to obtain more details.

2.4.4 Stackables Tools

The ERS stackables have numerous built-in tools that offer information on the health and well-being of the stack:

2.4.4.1 Stack Health Check

The stack health check feature provides a view into the overall operation of the stack; with information on the cascade connections, if the stack is resilient (return cable connected) or non-resilient, number of units in the stack and the model of each unit along with its unit number. This feature shows you a quick snapshot of the stack configuration and operation.

2.4.4.2 Stack Monitor

The stack monitor feature analyzes the health of a stack by monitoring the number of active units in the stack. With the stack monitor feature, when a stack is broken, the stack and any disconnected units from the stack send SNMP traps. If the stack or the disconnected units are still connected to the network, they generate log events and send trap messages to notify the administrator of the event. After the problem is detected, the stack and disconnected units continue to generate log events and send traps at a user-configurable interval until the situation is remedied (or the feature is disabled).

2.4.4.3 Stack Loopback Test

The stack loopback test feature allows the customer to quickly test the switch stack ports and the stack cables on the ERS units. This feature helps you while experiencing stack problems to determine whether the root cause is a bad stack cable or a damaged stack port and prevents potentially good switches being returned for service. You can achieve this by using two types of loopback tests – internal and external.

2.4.4.4 Stack Port Counters

The stack port counters show statistics of the traffic traversing the stacking connectors, including the size of packets, FCS errors, filtered frames, etc.

2.4.4.5 Environmental Information

This feature displays environmental information about the operation of the switch or units within a stack. The show environmental command does not require any configuration, and it reports the power supply status, fan status and switch system temperature.

2.4.4.6 CPU and Memory Utilization

The CPU utilization feature provides data for CPU and memory utilization. You can view CPU utilization information for the past 10 seconds, 1 minute, 1 hour, 24 hr, or since system boot up. The switch displays CPU utilization as a percentage. You can use CPU utilization information to see how the CPU is used during a specific time interval. The memory utilization provides you information on what percentage of the dynamic memory is currently used by the system. The switch displays memory utilization in terms of megabytes available since system boot up. This feature does not require a configuration. It is a display-only feature.

2.4.5 Key Health Indicator Collection

The VSP 9000 supports Key Health Indicators (KHI) that allow for the collection of statistics and information about the health of the system for troubleshooting purposes related to system failure. The Key Health Indicator (KHI) feature identifies a small number of key health indicators that allow quick assessment of the overall operational state of the VSP. These indicators do not provide complete coverage of all possible failure scenarios. Rather, KHI is a diagnostic tool for the health of the switch. Further debugging is required to correctly understand the system state and actions required to remedy the situation.

KHI provides global health information for the switch, including:

- Chassis health indication
- CPU performance health indication
- Port state change indication
- Forwarding health indication
- IP interface configuration and operation information
- Protocol information
- Management information: Log, TCP, UDP and Users

The switch stores the information locally and displays the information as requested by the user using show commands.

KHI supports multiple KHI types that track specific switch areas or subsystems. Each KHI type keeps track of the last ten events for the specific subsystem (for example, protocol going down or loss of connection) in a rolling history. KHI creates a reference point using a time stamp, and then tracks events from that point forward. Clear commands are provided to reestablish fresh timelines.

Generally, the KHI information allows you to track the source of a problem to a particular subsystem. Once this determination is made, you can use specific statistics for that subsystem (for example, OSPF-specific statistics and show commands) to further locate the source of the issue.

To configure KHI, you can enable or disable the feature globally. In addition, you can enable or disable some of the KHI types individually.

This additional control is provided for KHI types that have a greater impact on loaded systems.

The main configuration actions for KHI are:

- Enabling or disabling KHI (at global or feature-level)
- Displaying statistics
- Clearing statistics/history to establish a new timeline
- Currently, EDM does not support KHI configuration.

2.5 Security Features

The security of the Converged Campus is of paramount importance to the overall design. Avaya offers a variety of security mechanisms, both internally and externally, that work in conjunction to provide the highest level of security possible. This section focuses on the internal features built into the switches that guard against attacks on the network.

- Broadcast and Multicast Rate Limiting

  To protect the switch and other stations from a high number of broadcasts and multicasts, the switch has the ability to limit the broadcast/multicast rate. This feature can be configured on a per-port basis. By default, this feature is disabled, and should only be enabled in accordance with the recommendations from the previous section.

- Directed broadcast suppression

  The switches provide the ability to enable or disable forwarding for directed broadcast traffic on an IP-interface basis. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. When you disable or suppress directed broadcasts on an interface, all frames that are sent to the subnet broadcast address for a local router interface are dropped. Directed broadcast suppression protects hosts from possible Denial of Service (DoS) attacks.

- Prioritization of control traffic

  The switches use a very sophisticated prioritization scheme for scheduling received control packets (BPDUs, OSPF Hellos, etc.) on physical ports. This scheme involves two levels with both hardware and software queues and guarantees proper handling of these control packets regardless of the load on the switch. In turn, this guarantees the stability of the network. More specifically, it guarantees that the applications that heavily use broadcasts are handled with a lower priority. Note that you cannot use the CLI to view, configure, or modify these queues. Setting the queues and determining the type of packets entering each queue is Avaya confidential.

- ARP limitation

  The ARP request threshold limits the ability of the VSP 9000 to source ARP requests for workstation IP addresses it has not learned within its ARP table. The default setting for this function is 500 ARP requests per second. For networks experiencing excessive amounts of subnet scanning caused by a virus, Avaya recommends changing the ARP request threshold to a value between 100 and 50. This will help protect the CPU from causing excessive ARP requests, help protect the network, and lessen the spread of the virus to other PCs.

- Multicast Learning Limitation

  This feature protects the CPU from multicast data packet bursts generated by malicious applications such as viruses. Specifically, it protects against those viruses that cause the CPU to reach 100 percent utilization or that prevent the CPU from processing any protocol packets or management requests. If more than a certain number of multicast streams enter the CPU through a port during a sampling interval, the port is shut down until the user or administrator takes appropriate action.

- High Secure Mode

  To protect the VSP 9000 against IP packets with an illegal source address of 255.255.255.255 from being routed (per RFC 1812 Section 4.2.2.11 and RFC 971 Section 3.2), the VSP 9000 supports a configurable flag, called *high secure*. By default, this flag is disabled. Note that when you enable this flag, the desired behavior (not routing source packets with an IP address of 255.255.255.255) is applied.

- Access Policies

  Access policies let you control management access by setting policies for services to prevent or allow access to the switch. You can specify which hosts or networks can access the switch through FTP, http, rlogin/rsh, SSH, Telnet, and TFTP. You can set the access level (ro|rw|rwa).

- Configurable Software Daemons

  The ERS 8300 and VSP 9000 provide the ability to enable or disable various access methods. On the VSP 9000, you enable or disable ftp, tftp, telnet, rlogin, SSH, or SNMP. The ERS 8300 allows you to enable or disable ftp, tftp, telnet, rlogin, or SSH. The VSP 9000 also has a High Secure Mode in which all daemons are disabled; i.e., telnet, ftp, tftp, rlogin, and SNMP are disabled.

  For the ERS 5000, ERS 4500, ERS 2500, HTTP, telnet, and SNMP can be enabled or disabled through standard configuration.

- Route Policies

  The VSP 9000, ERS 8300, and ERS 5000 support IP RIP/OSPF accept/announce policies. This provides extra security by either blocking specific subnets or selecting to announce specific subnets.

- Port Lock Feature

  This feature lets you lock a port or prevent other users from changing port parameters or modifying port action. Locked ports cannot be modified in any way until the port is unlocked.

- Out-of-band Management

  Each Switch Fabric on the ERS 8300 and VSP 9000 provides an Out-of-band Management port. Traffic on this port is completely separated from the user traffic and provides a high secure network for management.

- Access Security

  The following access features are supported:

  - RADIUS authentication
  - TACACS+
  - SSH
  - SSL
  - Password security and CLI logging
  - SNMPv3

- Stopping IP Spoofed Packets

  Spoofed IP packets are stopped by configuring the VSP 9000 to ensure that IP packets are forwarded only if they contain the correct source IP address of your network. A spoofed packet is one that comes from the Internet into your network with a source address equal to one of the subnet addresses used on your network. Its source address belongs to one of the address blocks or subnets used on your network. With anti-spoofing protection, you have a filter rule/configuration assigned to the external interface, which examines the source address of all outside packets crossing that interface. If that address belongs to internal network or firewall itself, the packet is dropped. The correct source IP addresses consist of the IP network addresses that have been assigned to the site/domain. It is particularly important that you do this throughout your network, especially at the external connections to the existing Internet/upstream provider. By denying all invalid source IP addresses, you minimize the chances that your network will be the source of a spoofed DoS attack.

This will not stop DoS attacks coming from your network with valid source addresses, however. In order to prevent this, you need to know which IP network blocks are in use. You then create a generic filter that:

- Permits your sites' valid source addresses
- Denies all other source addresses

➢ Reverse Path Checking

The reverse path checking feature (available only on the VSP 9000), when enabled, prevents packet forwarding for incoming IP packets that have incorrect or forged (spoofed) IP addresses. Reverse path checking guarantees that traffic received on one interface was sent by a station from this interface, which prevents address spoofing. With this mode enabled, the VSP 9000 performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the packet is discarded.

You configure reverse path checking on a per-IP interface basis. When reverse path checking is enabled, the VSP 9000 checks all routing packets which come through that interface. It ensures that the source address and source interface appear in the routing table, and that it matches the interface on which the packet was received.

You can use one of two modes for reverse path checking:

- Exist-only mode: In this mode, reverse path checking checks whether the incoming packet's source IP address exists in the routing table. If the source IP entry is found, the packet is forwarded as usual; otherwise, the packet is discarded.
- Strict mode: In this mode, reverse path checking checks whether the incoming packet's source IP address exists in the routing table. If the source IP entry is not found, reverse path checking further checks if the source IP interface matches the packet's incoming interface. If they match, the packet is forwarded as usual; otherwise, the packet is discarded.

➢ DHCP Snooping

Dynamic Host Configuration Protocol (DHCP) snooping provides security to the network by preventing DHCP spoofing. DHCP spoofing refers to an attacker's ability to respond to DHCP requests with false IP information. DHCP snooping acts like a firewall between untrusted hosts and the DHCP servers, so that DHCP spoofing cannot occur. This feature classifies ports into trusted (where the DHCP server exists – which could simply be the uplink to the core) and untrusted (end user ports). DHCP requests are only forwarded to and thru the trusted ports. Any DHCP replies from untrusted ports are automatically dropped. A binding table is also created with the source MAC, IP address, VLAN, port, and lease time.

➢ Dynamic ARP Inspection

Dynamic Address Resolution Protocol (Dynamic ARP) inspection is a security feature that validates ARP packets in the network. Without dynamic ARP inspection, a malicious user can attack hosts, switches, and routers connected to the Layer 2 network by poisoning the ARP caches of systems connected to the subnet and by intercepting traffic intended for other hosts on the subnet. Dynamic ARP inspection prevents this type of *man-in-the-middle* attack. It intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings. The address binding table is dynamically built from information gathered in the DHCP request and reply when DHCP snooping is enabled. The MAC address from the DHCP request is paired with the IP address from the DHCP reply to create an entry in the DHCP binding table.

When Dynamic ARP inspection is enabled, IP traffic on *untrusted* ports is filtered based on the source MAC and IP addresses seen on the switch port. The switch forwards IP traffic when the source MAC and IP address matches an entry in the address binding table. Otherwise, the IP traffic is dropped. For dynamic ARP inspection to function, DHCP snooping must be globally enabled.

- Reverse Path Checking

  Reverse Path Checking is configured per IP interface. When Reverse Path Checking is enabled, the Ethernet Routing Switch 8000 checks all the routing packets that come through that interface, to ensure that the source address and source interface appear in the routing table and match the interface on the packet that was received. There are two modes for Reverse Path Checking:

  - Exist-only mode: checks whether the incoming packet's source IP address exists in routing table. If the source IP entry is found the packet is forwarded as normal; otherwise, the packet is discarded.
  - Strict mode: when configured in this mode, RPC first checks whether the incoming packet's source IP address exists in routing table. If the source IP entry is not found the packet is dropped; otherwise, RPC further checks if the source IP interface matches the packet's incoming interface. If they match, packet is forwarded as normal; otherwise, the packet is discarded.

  Exist-only mode should be configured for interfaces that accept traffic from other networks and participate in core networking IP transport. Strict mode should be configured for edge subnets where the only valid source IPs will be that of the IP interface itself.

  If you are unable to use Reverse Path Checking, create one or more filters for each interface, which examines the source address of all packets crossing that interface and rejects any spoofed packets.

- IP Source Guard

  IP Source Guard provides security to the network by filtering clients with invalid IP addresses. IP Source Guard is a Layer 2 (L2), port-to-port basis feature that works closely with information in the Dynamic Host Control Protocol (DHCP) snooping binding table. When you enable IP Source Guard on an untrusted port with DHCP snooping enabled, an IP filter entry is created or deleted for that port automatically, based on IP information stored in the corresponding DHCP snooping binding table entry. When a connecting client receives a valid IP address from the DHCP server, a filter is installed on the port to allow traffic only from the assigned IP address. A maximum of 10 IP addresses are allowed on each IP Source Guard-enabled port. When this number is reached, no more filters are set up and traffic is dropped.

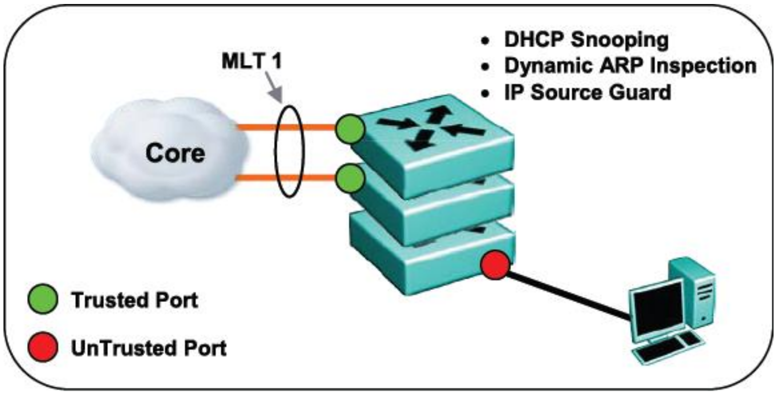

Figure 2.49: Edge Switch Security

- Untrusted ports are normally the end user access ports on the switch
- Trusted ports are normally the uplinks from the edge switch to the core
- Dynamic ARP Inspection and IP Source Guard use the binding table created by DHCP Snooping, therefore, in order to use these features, DHCP Snooping must be enabled.
- IP Source Guard should not be enabled on uplink ports from the Edge to the Core, only enable on Edge access ports (untrusted ports) where DHCP Snooping and Dynamic ARP Inspection are enabled.

2.6 Network Management

For the Super Large Campus, a full suite of management tools is typically requested because there are usually a large number of devices that need to be managed. The key is to simplify the management process while still supporting a large number of devices and potentially complex network designs.

2.6.1 Access Security

Ensuring secure access to the infrastructure is extremely important for maintaining the integrity of the Super Large Campus network. There are a few basic guidelines as part of the Avaya best practices as highlighted here.

- SSH (secure shell) provides encrypted communication to the Ethernet Switches, this is preferable over Telnet which is clear text
- A Web interface is available for the VSP 9000 and the ERS products. If you use this interface, Avaya strongly recommends implementing SSL and/or change the port number to which the switch will be managed via the Web.
- Create access lists to ensure only certain clients or subnets are able to manage the switches.
- Implement RADIUS for authorization and accounting.
  - RADIUS fallback ensures access to switches in the event the RADIUS server is not reachable.
  - Change all local passwords from factory default.
- TACACS+ is also supported for authorization and accounting.
- Configure remote syslog capabilities to ensure all critical log file information is kept and recorded off the switches.
- Configure SNMPv3 for secure SNMP communications to all Ethernet Switches.
  - Make sure community strings are changed from default.

2.6.2 Network Management

Avaya provides a comprehensive set of solutions and tools to enable a system-wide life cycle management. The objective is to offer management solutions that are efficient, survivable, consistent, and drive simplicity. This is accomplished by addressing all areas of FCAPS.

FCAPS is the ISO Telecommunications Management Network model and framework for network management. FCAPS is an acronym for *Fault, Configuration, Accounting, Performance* and *Security* which are the management categories into which the ISO model defines network management tasks.

The goal of fault management is to recognize, isolate, correct and log faults that occur in the network. It uses trend analysis to predict errors so that the network is always available. This is established by monitoring different aspects of the network for abnormal behavior.

Configuration Management covers configuration changes to the network, including adding new devices, removing old ones and modifying the configuration of others. It ensures that changes are conducted consistently in an ordered manner no matter how many changes are required.

Accounting Management gathers usage statistics, and by using the statistics, the users can be billed and usage quota enforced. For non-billed networks, "administration" replaces "accounting". The goals of administration are to administer the set of authorized users by establishing passwords, permissions, and to administer network operations.

Performance Management addresses the throughput, percentage utilization, error rates and response times of the network. By collecting and analyzing performance data, the network health is monitored. Trends indicate capacity or reliability issues before they become service affecting. Performance thresholds can trigger alarms to proactively correct potential bandwidth issues.

Security Management is the process of controlling access to assets in the network. Data security can be achieved mainly with authentication i.e. verifying the identity of the person accessing the network, authorization (that is determining what the user can than do) and finally encryption (that is protecting traffic from unauthorized interception when it passes across the network).

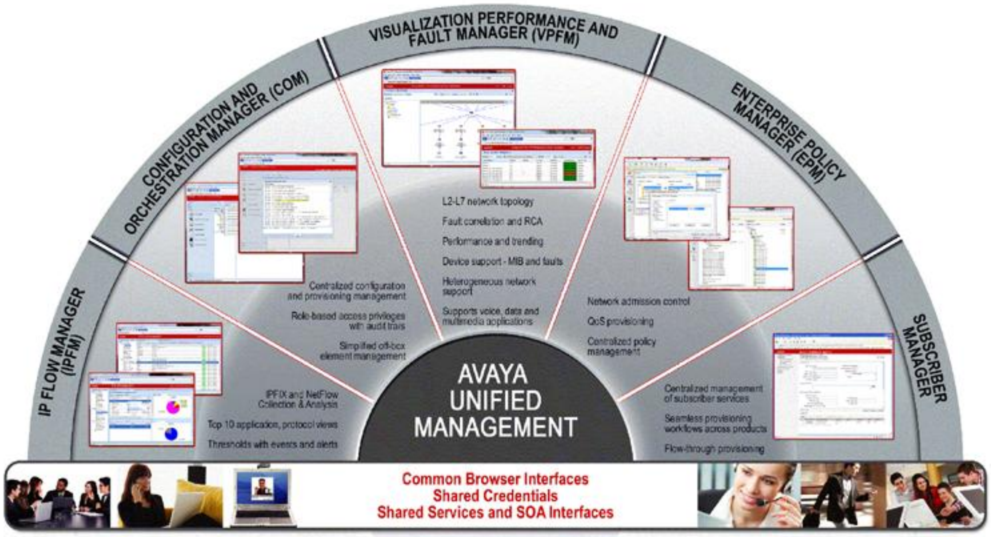
IP FLOW MANAGER (IPFM)
CONFIGURATION AND ORCHESTRATION MANAGER (COM)
VISUALIZATION PERFORMANCE AND FAULT MANAGER (VPFM)
ENTERPRISE POLICY MANAGER (EPM)
SUBSCRIBER MANAGER
IPFIX and NetFlow Collection & Analysis
Top 10 application, protocol views
Thresholds with events and alerts
Centralized configuration and provisioning management
Role-based access privileges with audit trails
Simplified off-box element management
L2-L7 network topology
Fault correlation and RCA
Performance and trending
Device support - MIB and faults
Heterogeneous network support
Supports voice, data and multimedia applications
Network admission control
QoS provisioning
Centralized policy management
Centralized management of subscriber services
Seamless provisioning workflows across products
Flow-through provisioning
AVAYA UNIFIED MANAGEMENT
Common Browser Interfaces
Shared Credentials
Shared Services and SOA Interfaces

Figure 2.50: Unified Communications Management

2.6.3 Unified Communications Management (UCM)

In essence, Avaya Unified Communications Management is the 'glue' that binds together distributed network management components to establish a true unified communications environment. UCM, which follows the FCAPS framework, provides comprehensive and simplified network management across voice, data and multimedia applications through integrating management features and capabilities including visualization, configuration, fault and performance management, subscriber services and security.

UCM is a centralized and integrated set of management tools and applications. The key and distinguishing element of UCM is the in-built Common Services that are an integral part of each of the different network management applications that form UCM. Common Services allow network management applications to integrate with each other so that common components (for example, user data, database information and certificate management) can be shared, without requiring that the same definitions and configurations be repeated for each application. The result is a fully integrated, single point of contact, providing a unified view of the network, while streamlining workflows and the management environment, and reducing installation and configuration time, system operations and maintenance.

Each product offering in the Avaya UCM solution has two components; first, the application which provides the application-specific functions by purpose. These components include:

- ➢ **Visualization Performance and Fault Manager** (data & voice)
  - ▪ Multi-vendor network discovery, root cause analysis, network topology maps

- **Configuration and Orchestration Manager** (data)
    - Multi-element configuration for Ethernet Routing Switches
- **Enterprise Policy Manager (data)**
    - Network access control policies, bandwidth management, QoS
- **IP Flow Manager** (data)
    - IPFIX collection, analysis and reporting
- **Proactive Voice Quality Management** (voice)
    - Centralized management of telephony subscriber services

Second is Common Services (UCM-CS). All applications in the UCM solution have the same UCM-CS which offer a common set of tools, interfaces, service subscriptions, user authentication, and user definition / management. Integration between components occurs at the UCM-CS layer.

The UCM key features include:

- Application co-residency – lower CAPEX/OPEX
- Single unified management domain – decreased complexity
- Integrated workflow – reduced errors
- Centralized authentication and navigation – improved user experience
- Shared data – reduced errors, secure
- Centralized management via browser – simple to use
- Single Sign-On – simple to use, secure
- Simplified system admin configuration – simple to use
- Flexible XML Architecture – investment protection

2.6.4 Visualization Performance & Fault Manager

Visualization Performance & Fault Manager (VPFM) – part of the Unified Communications Management (UCM) solution – provides a heterogeneous network discovery and visualization tool that uses standards-based and proprietary protocols to provide an end-to-end view of the network. The functionality available with VPFM includes:

- Network Device and End-node Discovery
    - Proprietary-based (Avaya SONMP, Cisco CDP)
    - Standards-based (IP, SNMP, IEEE 802.1AB, ad etc.)
    - Application Discovery
    - Servers, PC, etc.
- L2 and L3 Network Topology Visualization
    - Device relationships and their interconnectivity
    - Switch views
    - Campus Views
    - VoIP Service (Application) Views
- Fault Management

  - Event Correlation and Analysis
  - Event Handling and Scripting
  - Trap Retention and Exporting
  - Custom Propagation Rules
- Device Performance Monitoring
  - LAG Performance Monitoring
  - Threshold support
  - Trending, Graphing, and data Exporting
  - Custom Device Classification (scope) Management
- Enhanced Diagnostics
  - L2 and L3 Diagnostic management

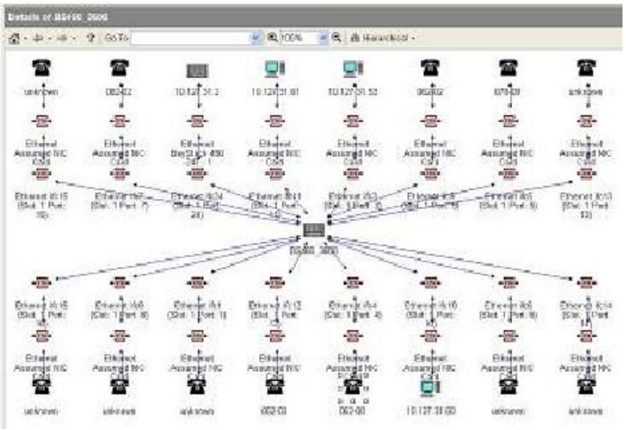

Figure 2.51: VPFM Topology View

2.6.5 Configuration & Orchestration Manager

Configuration & Orchestration Manager (COM) delivers centralized configuration, provisioning and troubleshooting for enterprise-wide devices and technologies. It is a real-time Web-based application, enabling anytime, anywhere access for multiple authorized users. With role-based access privileges and audit trails, a high level of security is ensured.

COM simplifies network-wide configuration by offering a topology-driven GUI with workflow-based visibility. Complex configurations such as VLAN, SMLT/MLT, VRF and multicast are easily managed using templates and wizards.

In addition, COM can provide network discovery, password management and serve as a platform for downloading, launching and hosting off-box device management. With centralized off-box element management, COM also streamlines the managing of inventory, updates as well as device and user access controls.

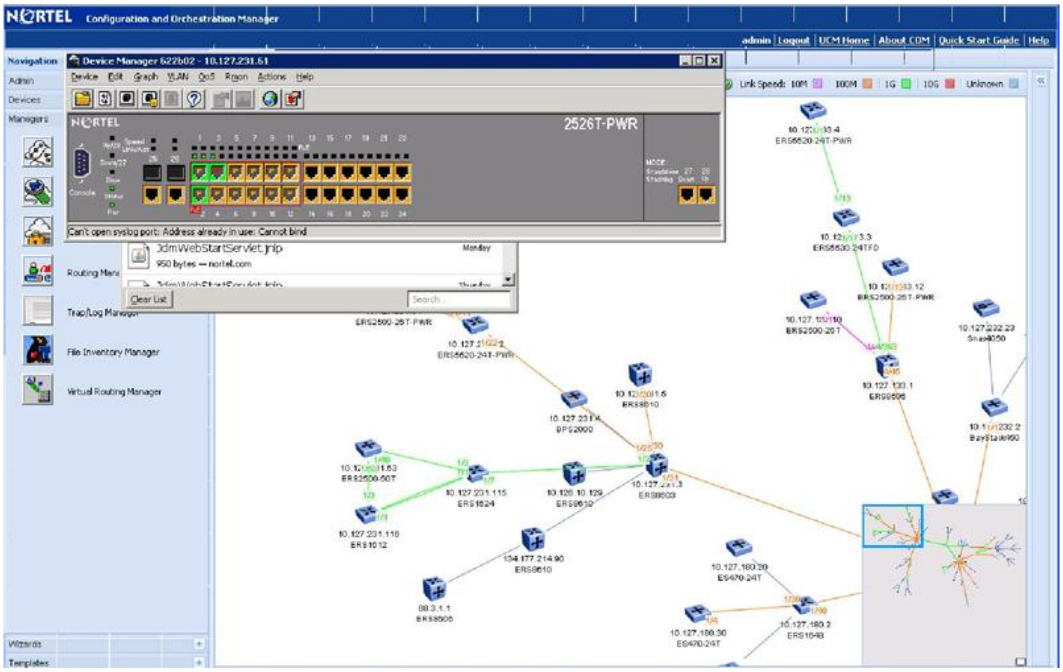

Can't open syslog port: Address already in use: Cannot bind

Figure 2.52: COM Topology View

COM also provides a launch point for Enterprise Device Manager (EDM) and off-box EDM functionality. EDM is the on-box management platform for all Ethernet switching products. EDM replaces Java Device Manager (JDM) and delivers a more streamlined approach to single element management. EDM uses a web-based interface to simplify configuration of the network element.

EDM comes as part of the software image for the Ethernet Routing Switches at no additional cost. The look and feel of EDM is very similar to that of legacy JDM, so transitioning to the new management platform should be nearly seamless. Because EDM is an on-box configuration tool, there are certain features that cannot be done on-box (syslog, trap receiver, audit logs, mib browser, etc.) and therefore requires an off-box application – COM provides this off-box application through EDM plug-ins. With COM and EDM, the network administrator has a full suite of tools required to quickly and easily perform configuration and orchestration of the network infrastructure.

2.6.6 Enterprise Policy Manager

Avaya Enterprise Policy Manager (EPM) is a network-level application that allows administrators to manage network bandwidth, prioritize traffic streams, and set network access policies. EPM enables critical applications to receive the proper QoS by deploying common policies from one central management application. This ensures consistency across the network along with simplifying the deployment of an enterprise-wide QoS policy.

EPM can be used to define Policy Enforcement Points in the network. An operator will associate multiple ports throughout the network across many switches to a Role within EPM. A user can then quickly define a Policy (consisting of a Traffic Condition, Action, and Schedule) into a role which will once applied immediately apply that policy to all of the switches/ports in that Role.

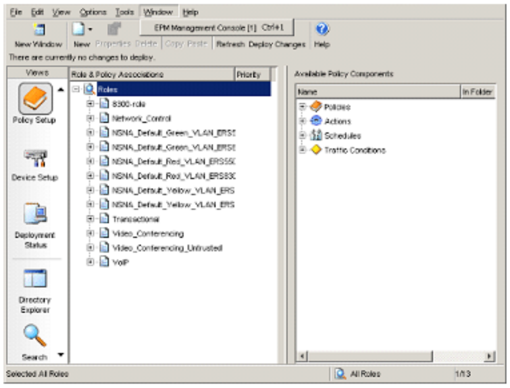

There are currently no changes to deploy.

- Roles
  - 8300-role
  - Network_Control
  - NSNA_Default_Green_VLAN_ERS
  - NSNA_Default_Green_VLAN_ERS
  - NSNA_Default_Red_VLAN_ERS55
  - NSNA_Default_Red_VLAN_ERS83
  - NSNA_Default_Yellow_VLAN_ERS
  - NSNA_Default_Yellow_VLAN_ERS
  - Transactional
  - Video_Conferencing
  - Video_Conferencing_Untrusted
  - VoIP

- Policies
- Actions
- Schedules
- Traffic Conditions

Figure 2.53: Enterprise Policy Manager

2.6.7 IP Flow Manager

IP Flow Manager (IPFM) is the external IPFIX collector used to analyze the various services and applications on the network. IPFIX is a very powerful tool for collecting application flows across the network. This enables the network administrator to gain insight in regard to usage monitoring of the network, user activity, network abuse, top users and top protocols, and can help answer the questions about true end-to-end application performance.

Note – Ethernet Routing Switches support IP Flow Manager and IPFIX. However, the VSP 9000 does not currently support these features.

An IP flow is defined as a set of packets sent over a period of time that have some common properties. These properties include:

- Source IP address
- Destination IP address
- Protocol type
- Source protocol port
- Destination protocol port
- Ingress VLAN ID
- Ingress port and observation point (VLAN or port)

IPFM turns IPFIX data into useful and easily understood information. This valuable information now increases the visibility into whom and which applications are consuming network resources and bandwidth. It also is used as a capacity planning tool, allowing the network administrator to proactively make changes to avert any potential bandwidth issues that may arise with new applications being rolled out continually.

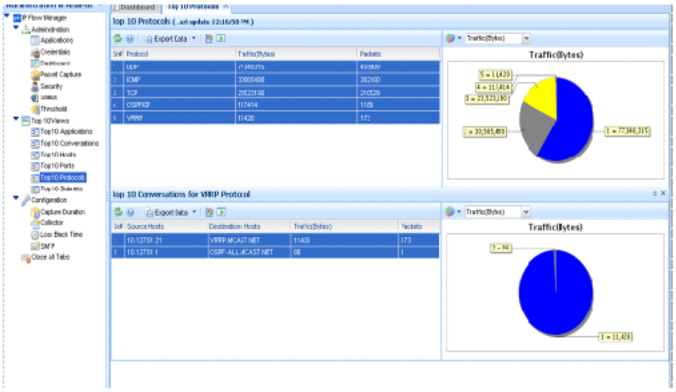

Top 10 Protocols (Last update 12:16:50 PM)

Sr#	Protocol	Traffic(Bytes)	Packets
1	UDP	[illegible]	[illegible]
2	ICMP	30565490	302960
3	TCP	23523190	210529
4	OSPFIGP	117414	1169
5	VRRP	11420	172

Top 10 Conversations for VRRP Protocol

Sr#	Source Hosts	Destination Hosts	Traffic(Bytes)	Packets
1	10.127.51.21	VRRP.MCAST.NET	11420	173
1	10.127.51.1	OSPF-ALL.MCAST.NET	96	1

Figure 2.54: IP Flow Manager

2.6.8 Proactive Voice Quality Management

Business challenge – Unified communications can give your business a competitive advantage by empowering the end user with productivity-improving applications, but deploying and maintaining Voice over IP (VoIP) can be challenging. The problem facing many customers is summarized in the benchmark survey Benchmarking VoIP Performance Management by the Aberdeen Group (March 2008). “Since transitioning from old-fashioned dedicated analog telephony to VoIP is a major initiative at many organizations, managing performance of VoIP services is becoming increasingly important. Managing the performance of voice services delivered over IP networks is possible through execution of a variety of strategies; but the challenge that organizations face is in identifying the right strategy and developing capabilities needed for successful execution of that strategy.” How do you know that your network is ready to handle the real-time bandwidth required? How can you monitor the end-user experience versus the network performance? How will you troubleshoot problems in real-time as they occur and reduce the number of support requests?

The Solution – PVQM (Proactive Voice Quality Management). Co-developed by Avaya and NetIQ, Proactive Voice Quality Management ensures that you have the right tools in place to support the right processes to stay on top of your VoIP service quality. Count on Avaya and NetIQ to provide your VoIP Service Level Management needs.

What is Proactive Voice Quality Management? Proactive Voice Quality Management offers a life-cycle approach that provides the necessary management tools to support each phase of an IP Telephony project: assessment, pre-deployment, ongoing monitoring and reporting. PVQM then provides a set of interrelated technologies that enable an effective Service Level Management solution for IP Telephony. PVQM focuses on the end-user Quality of Experience (QoE) using standards-based technologies ensuring you can stay on top of your service quality needs.

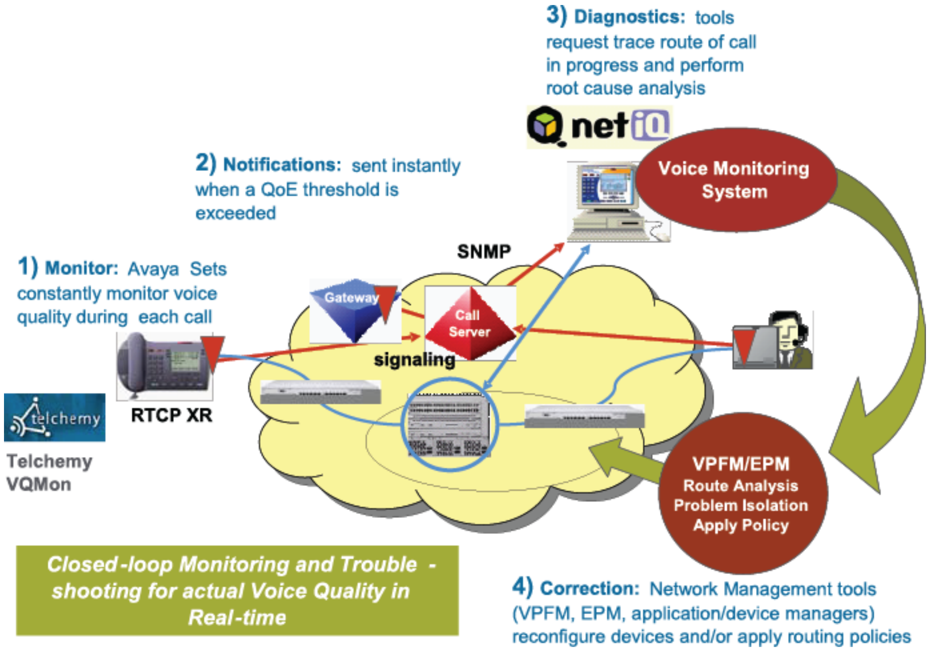
3) Diagnostics: tools request trace route of call in progress and perform root cause analysis
netiQ
Voice Monitoring System
2) Notifications: sent instantly when a QoE threshold is exceeded
SNMP
1) Monitor: Avaya Sets constantly monitor voice quality during each call
Gateway
Call Server
signaling
RTCP XR
telchemy
Telchemy VQMon
VPFM/EPM
Route Analysis
Problem Isolation
Apply Policy
Closed-loop Monitoring and Trouble-shooting for actual Voice Quality in Real-time
4) Correction: Network Management tools (VPFM, EPM, application/device managers) reconfigure devices and/or apply routing policies

Figure 2.55: PVQM

Network assessment and pre-deployment – Is your network ready for VoIP? According to a Gartner Group study, 85 percent of locations will require upgrades to their data networks to properly run voice. Before deploying VoIP, you need to have the proper equipment and identify if you need network upgrades before deployment. Fully supporting Avaya Network Health Checks, NetIQ's Vivinet Assessor determines quickly and easily how well VoIP will work on the network prior to deployment. Before you invest in costly training and pilot deployments, Vivinet Assessor predicts the overall call quality that you can expect from the network and generates polished, customizable reports detailing the network's VoIP readiness. Also in this phase, Enterprise Policy Manager provides centralized management of your network QoS through an intuitive graphical user interface. Deploy QoS in a consistent manner across your enterprise network using Enterprise Policy Manager.

Monitoring and reporting - Managing VoIP availability and quality needs – Once VoIP is deployed, you must monitor the VoIP environment to maximize uptime, reliability and quality. NetIQ's AppManager Suite is a robust platform and suite of modules that provide comprehensive monitoring, management and reporting for voice solutions. NetIQ AppManager for Avaya extends the suite to optimize the health and availability of Avaya's VoIP platforms. Call quality is monitored from an end-user perspective and in real time with embedded software from Telchemy. Telchemy's VQmon is integrated into Avaya's IP phones to provide the metrics needed to support management and QoS reporting protocols and facilitate problem diagnosis. NetIQ AppManager hen combines platform, system and call quality metrics for Avaya Call Servers with information about the availability and health of network devices and overall network performance for VoIP. Avaya's Visualization Performance Faults Manager can also be deployed providing integrated infrastructure management for your data network, data services and IP Telephony. With Avaya Communication Server 1000 (CS 1000), VoIP quality is monitored from an end-user perspective providing a high degree of correlation between how network performance affects actual service quality. End-user QoE is monitored and reported in real-time using industry-standard technology (ITU G.107 and IETF 3611) providing the first open solution on the market. This allows operators to focus on real problems in the network. Furthermore, critical insight is provided to highlight the underlying conditions that led to service quality degradation. This allows operators to map VoIP service quality back to the underlying network infrastructure once and for all!

PVQM Summary – Successful unified communications implementation with VoIP requires integrated management solutions that allow you to take control of your entire voice network and server infrastructure. Understanding how data traffic will affect voice applications — before deployment — and then continually monitoring and diagnosing the status of IP Telephony devices will help maximize success. For all stages of deployment, NetIQ's products provide the most comprehensive solution available on the market. NetIQ delivers assessment, monitoring and diagnostic products to help you ensure a successful implementation and accelerate your overall return on VoIP investment. In addition, Enterprise Policy Manager and Visualization Performance Faults Manager complement your investment in Avaya data solutions, providing point-and-click management of your converged network.

3. Configuration Example

The following example will be used throughout the remainder of this document as the reference topology for configuration of all the best practice recommended parameters. This configuration example only provides a fraction of the capabilities supported by Avaya switches. Additional security and authentication parameters can be optionally enabled and are covered by additional configuration guides available for download on the Avaya technical support site http://support.avaya.com.

Configuration Design Details:

VLAN Name	VLAN ID	IP Subnet
IST - Core	2	1.1.1.0/30
IST – Edge Distribution	2	1.1.1.4/30
IST – Data Center Distribution	2	1.1.1.8/30
Edge-SMLT-2	4	10.0.4.0/24
DC-SMLT-3	5	10.0.5.0/24
Management	2000	10.20.0.0/24
DataCenter	2010	10.20.10.0/24
User1	2101	10.21.1.0/24
Voice1	2102	10.21.2.0/24
User2	2201	10.22.1.0/24
Voice2	2202	10.22.2.0/24

Table 3.0 – Configuration Details

Management VLAN

- Configured as Layer 2 across the network

Layer 3 Edge Routing

In order to provide examples of both RSMLT Layer 2 Edge and VRRP, the following configurations will be used:

- User and Voice VLANs are using VRRP / Backup Master
- Data Center VLANs are using RSMLT Layer 2 Edge

Figure 3.1 shows the entire test topology that was used in the validation of the Super Large Campus design. The configuration examples that follow are from the top shaded portion of the diagram. The bottom portion's configuration would be identical to the top with the only difference being VLAN ID's and IP addresses.

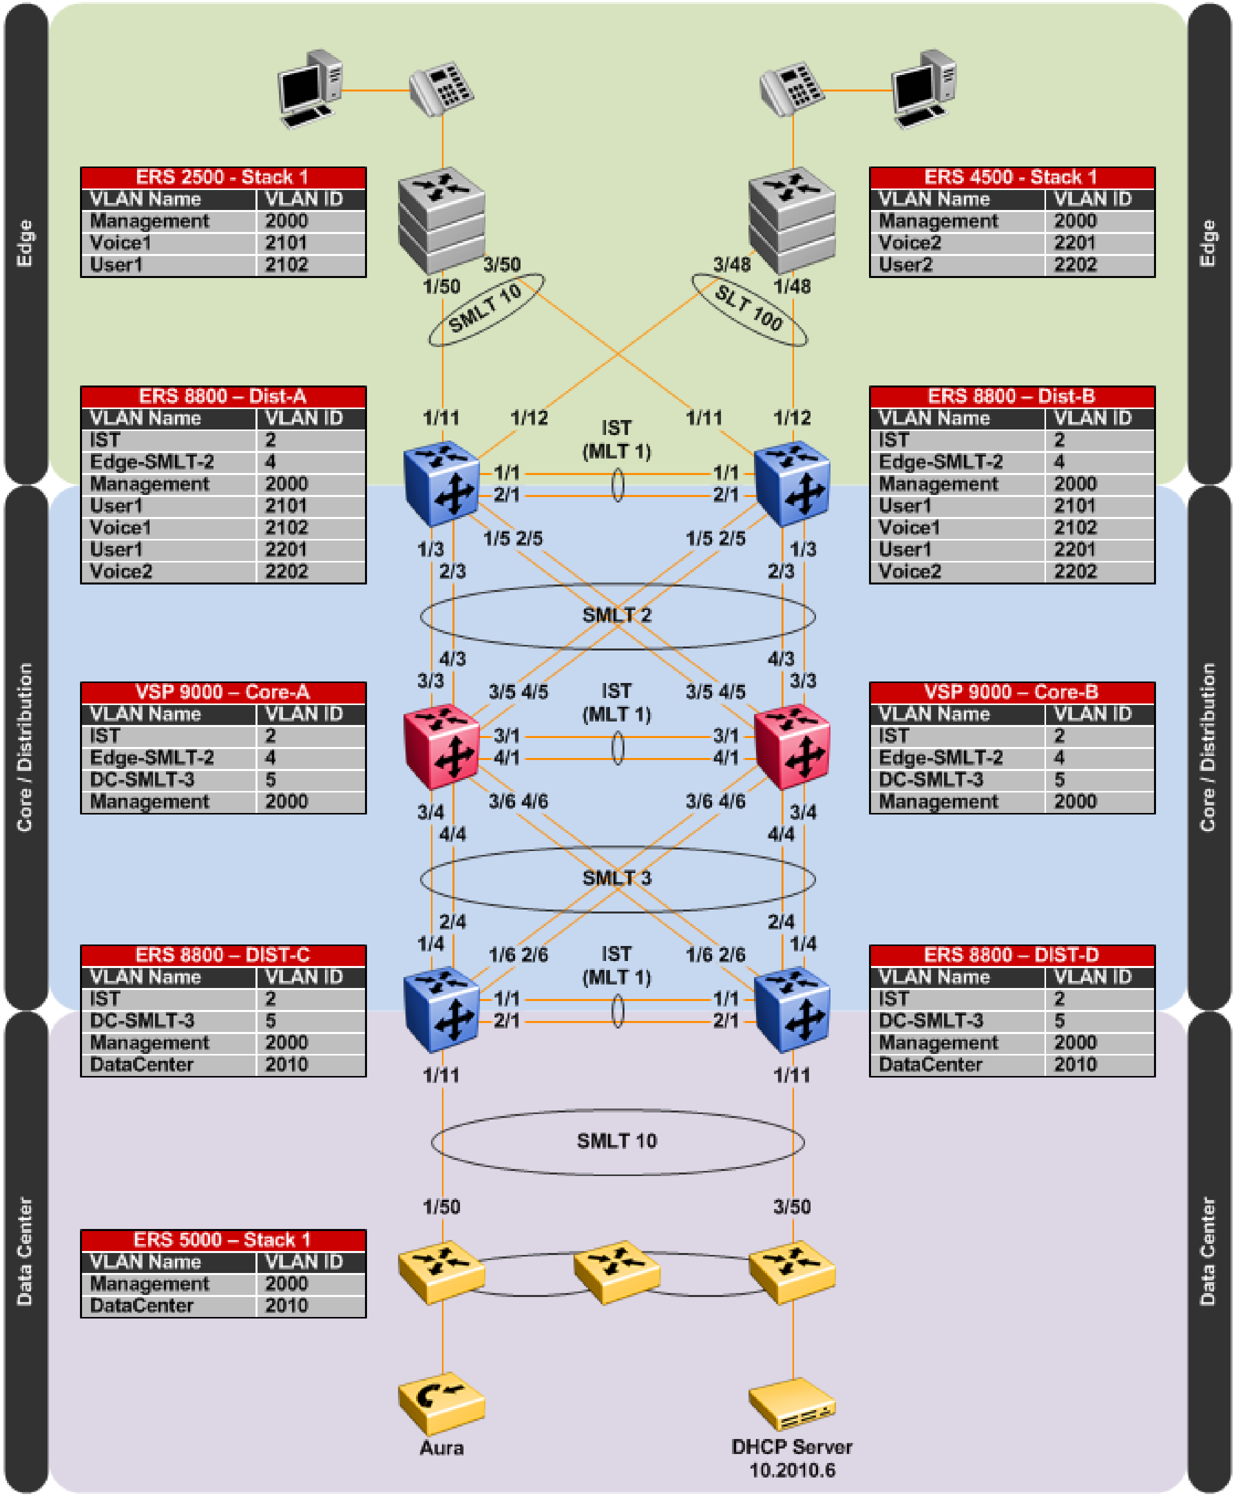

ERS 2500 - Stack 1	
VLAN Name	VLAN ID
Management	2000
Voice1	2101
User1	2102

ERS 4500 - Stack 1	
VLAN Name	VLAN ID
Management	2000
Voice2	2201
User2	2202

ERS 8800 – Dist-A	
VLAN Name	VLAN ID
IST	2
Edge-SMLT-2	4
Management	2000
User1	2101
Voice1	2102
User1	2201
Voice2	2202

ERS 8800 – Dist-B	
VLAN Name	VLAN ID
IST	2
Edge-SMLT-2	4
Management	2000
User1	2101
Voice1	2102
User1	2201
Voice2	2202

VSP 9000 – Core-A	
VLAN Name	VLAN ID
IST	2
Edge-SMLT-2	4
DC-SMLT-3	5
Management	2000

VSP 9000 – Core-B	
VLAN Name	VLAN ID
IST	2
Edge-SMLT-2	4
DC-SMLT-3	5
Management	2000

ERS 8800 – DIST-C	
VLAN Name	VLAN ID
IST	2
DC-SMLT-3	5
Management	2000
DataCenter	2010

ERS 8800 – DIST-D	
VLAN Name	VLAN ID
IST	2
DC-SMLT-3	5
Management	2000
DataCenter	2010

ERS 5000 – Stack 1	
VLAN Name	VLAN ID
Management	2000
DataCenter	2010

Figure 3.1: Configuration Topology

3.1 Software Versions & Upgrade Policy

The Super Large Campus Solution was validated with specific software versions as indicated below. These software releases provided stable interoperability of products at the time of release of Super Large Campus Solution. However, post release, individual products will undergo minor updates or major release upgrades. This section provides guidance on how to benefit from periodic software changes while still maintaining the overall integrity of the solution.

In general, upgrading of products to the next major release moves the solution outside the boundaries of validation undertaken for the particular solution. Major product upgrades will be validated in future updated versions of the Super Large Campus Solution Guide. If upgrades to the next major revisions are required, consult with Avaya and proceed with caution. Review the Release Notes carefully to ensure changes in the newly released software will not adversely affect the overall network.

In contrast to the above, you can update or patch products within the Super Large Campus Solution to the next minor software release with the worry of moving the solution outside the validation boundaries. This allows you to take advantage of bug fixes as they become available.

Software versions used for the Super Large Campus validation:

- VSP 9000 Release 3.0
- ERS 8800 Release 7.0
- ERS 8300 Release 4.2.2.2
- ERS 5000 Release 6.2
- ERS 4500 Release 5.4
- ERS 2500 Release 4.3

Minor software releases for the above products are incremented as required for each maintenance release. When downloading software from the Avaya support site, each version of software is clearly identified by type (Major or Maintenance).

3.2 Core Switch Configuration

This section covers the basic configuration of the Virtual Services Platform 9000 switch cluster core.

Core-A
VLAN ID:2 (IST) IP: 1.1.1.1/30
VLAN ID: 4 (Edge-SMLT-2) IP: 10.0.4.1/24
VLAN ID: 5 (DC-SMLT-3) IP: 10.0.5.1/24
VLAN ID: 2000 (Management) IP: 10.20.0.1/24
Loopback ID: 1 IP: 1.0.0.1/32
MLT ID: 1 (IST) VLAN IDs: 2,4,5,2000 Ports: 3/1,4/1
MLT ID: 2 (Edge-SMLT-2) VLAN IDs: 4,2000 Ports: 3/3,3/5,4/3,4/5
MLT ID: 3 (DC-SMLT-3) VLAN IDs: 5,2000 Ports: 3/4,3/6,4/4,4/6

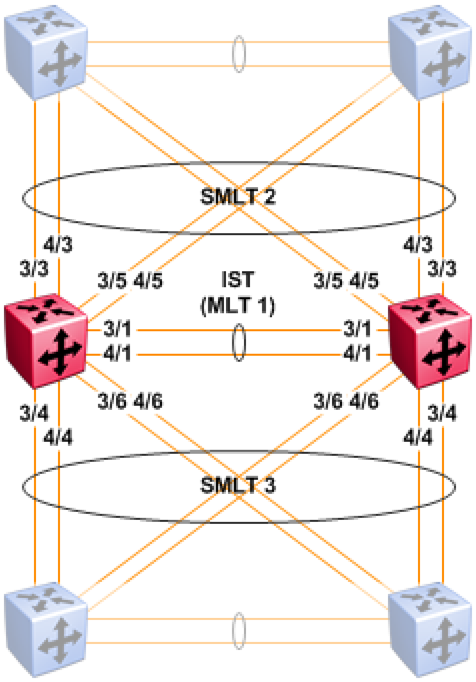

Core-B
VLAN ID:2 (IST) IP: 1.1.1.2/30
VLAN ID: 4 (Edge-SMLT-2) IP: 10.0.4.2/24
VLAN ID: 5 (DC-SMLT-3) IP: 10.0.5.2/24
VLAN ID: 2000 (Management) IP: 10.20.0.1/24
Loopback ID: 1 IP: 1.0.0.2/32
MLT ID: 1 (IST) VLAN IDs: 2,4,5,2000 Ports: 3/1,4/1
MLT ID: 2 (Edge-SMLT-2) VLAN IDs: 4,2000 Ports: 3/3,3/5,4/3,4/5
MLT ID: 3 (DC-SMLT-3) VLAN IDs: 5,2000 Ports: 3/4,3/6,4/4,4/6

Figure 3.2 – Core Switch Cluster Configuration

ⓘ A second Core SMLT VLAN may optionally be deployed between the VSP 9000 cluster core and the ERS 8800 distribution switches to improve OSPF recovery times if the OSPF designated router fails.

3.2.1 Virtual LANs

For this configuration step the following VLAN and IP parameters will be enabled:

VSP 9000 – Core-A		
VLAN Name	**VLAN ID**	**IP Address**
IST	2	1.1.1.1/30
Edge-SMLT-2	4	10.0.4.1/24
DC-SMLT-3	5	10.0.5.1/24
Management	2000	10.20.0.1/24
VSP 9000 – Core-B		
VLAN Name	**VLAN ID**	**IP Address**
IST	2	1.1.1.2/30
Edge-SMLT-2	4	10.0.4.2/24
DC-SMLT-3	5	10.0.5.2/24
Management	2000	10.20.0.2/24

Table 3.2.1 – Core-A / Core-B VLAN and IP Interface Parameters

1 Create IST, SMLT and Management VLANs:

```
Core-A:1(config)# vlan create 2 name "IST" type port-mstprstp 1
Core-A:1(config)# vlan mac-address-entry 2 aging-time 21601
Core-A:1(config)# vlan create 4 name "Edge-SMLT-2" type port-mstprstp 1
Core-A:1(config)# vlan mac-address-entry 4 aging-time 21601
Core-A:1(config)# vlan create 5 name "DC-SMLT-3" type port-mstprstp 1
Core-A:1(config)# vlan mac-address-entry 5 aging-time 21601
Core-A:1(config)# vlan create 2000 name "Management" type port-mstprstp 1
```

```
Core-B:1(config)# vlan create 2 name "IST" type port-mstprstp 1
Core-B:1(config)# vlan mac-address-entry 2 aging-time 21601
Core-B:1(config)# vlan create 4 name "Edge-SMLT-2" type port-mstprstp 1
Core-B:1(config)# vlan mac-address-entry 4 aging-time 21601
Core-B:1(config)# vlan create 5 name "DC-SMLT-3" type port-mstprstp 1
Core-B:1(config)# vlan mac-address-entry 5 aging-time 21601
Core-B:1(config)# vlan create 2000 name "Management" type port-mstprstp 1
```

2 Add IP Addresses to the IST and SMLT VLANs:

```
Core-A:1(config)# interface vlan 2
Core-A:1(config-if)# ip address 1.1.1.1 255.255.255.252
Core-A:1(config-if)# exit
Core-A:1(config)# interface vlan 4
Core-A:1(config-if)# ip address 10.0.4.1 255.255.255.0
Core-A:1(config-if)# exit
Core-A:1(config)# interface vlan 5
Core-A:1(config-if)# ip address 10.0.5.1 255.255.255.0
Core-A:1(config-if)# exit
Core-A:1(config)# interface vlan 2000
Core-A:1(config-if)# ip address 10.20.0.1 255.255.255.0
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 2
Core-B:1(config-if)# ip address 1.1.1.2 255.255.255.252
Core-B:1(config-if)# exit
Core-B:1(config)# interface vlan 4
Core-B:1(config-if)# ip address 10.0.4.2 255.255.255.0
Core-B:1(config-if)# exit
Core-B:1(config)# interface vlan 5
Core-B:1(config-if)# ip address 10.0.5.2 255.255.255.0
Core-B:1(config-if)# exit
Core-B:1(config)# interface vlan 2000
Core-B:1(config-if)# ip address 10.20.0.2 255.255.255.0
Core-B:1(config-if)# exit
```

3.2.2 Switch Clustering

For this configuration step the following switch clustering parameters will be enabled:

VSP 9000 – Core-A and Core-B		
MLT ID	**Port(s)**	**VLAN Members**
1 (IST)	3/1,4/1	2,4,5,2000
2 (Edge-SMLT-2)	3/3,3/5,4/3,4/5	4,2000
3 (DC-SMLT-3)	3/4,3/6,4/4,4/6	5,2000

Table 3.2.2 – Core-A / Core-B Switch Clustering Parameters

1 Create MLT 1 for IST and assign ports:

```
Core-A:1(config)# mlt 1 enable name “IST-Core-B”
Core-A:1(config)# mlt 1 member 3/1,4/1
Core-A:1(config)# mlt 1 encapsulation dot1q
Core-A:1(config)# vlan mlt 2 1
```

```
Core-B:1(config)# mlt 1 enable name “IST-Core-A”
Core-B:1(config)# mlt 1 member 3/1,4/1
Core-B:1(config)# mlt 1 encapsulation dot1q
Core-B:1(config)# vlan mlt 2 1
```

2 Create IST:

```
Core-A:1(config)# interface mlt 1
Core-A:1(config-if)# ist peer-ip 1.1.1.2 vlan 2
Core-A:1(config-if)# ist enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface mlt 1
Core-B:1(config-if)# ist peer-ip 1.1.1.1 vlan 2
Core-B:1(config-if)# ist enable
Core-B:1(config-if)# exit
```

3 Create SMLT 2 to ERS 8800 Dist-A and Dist-B Switches:

```
Core-A:1(config)# mlt 2 enable name “SMLT-Dist-AB”
Core-A:1(config)# mlt 2 member 3/3,3/5,4/3,4/5
Core-A:1(config)# mlt 2 encapsulation dot1q
Core-A:1(config-if)# interface mlt 2
Core-A:1(config-if)# smlt
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# mlt 2 enable name “SMLT-Dist-AB”
Core-B:1(config)# mlt 2 member 3/3,3/5,4/3,4/5
Core-B:1(config)# mlt 2 encapsulation dot1q
Core-B:1(config-if)# interface mlt 2
Core-B:1(config-if)# smlt
Core-B:1(config-if)# exit
```

4 Create SMLT 3 to ERS 8800 Dist-C and Dist-D Switches:

```
Core-A:1(config)# mlt 3 enable name "SMLT-Dist-CD"
Core-A:1(config)# mlt 3 member 3/4,3/6,4/4,4/6
Core-A:1(config)# mlt 3 encapsulation dot1q
Core-A:1(config-if)# interface mlt 3
Core-A:1(config-if)# smlt
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# mlt 3 enable name "SMLT-Dist-CD"
Core-B:1(config)# mlt 3 member 3/4,3/6,4/4,4/6
Core-B:1(config)# mlt 3 encapsulation dot1q
Core-B:1(config-if)# interface mlt 3
Core-B:1(config-if)# smlt
Core-B:1(config-if)# exit
```

5 Add VLANs 4, 5 and 2000 to the IST:

```
Core-A:1(config)# vlan mlt 4 1
Core-A:1(config)# vlan mlt 5 1
Core-A:1(config)# vlan mlt 2000 1
```

```
Core-A:1(config)# vlan mlt 4 1
Core-A:1(config)# vlan mlt 5 1
Core-A:1(config)# vlan mlt 2000 1
```

6 Add VLANs 4 and 2000 to SMLT 2:

```
Core-A:1(config)# vlan mlt 4 2
Core-A:1(config)# vlan mlt 2000 2
```

```
Core-A:1(config)# vlan mlt 4 2
Core-A:1(config)# vlan mlt 2000 2
```

7 Add VLANs 5 and 2000 to SMLT 3:

```
Core-A:1(config)# vlan mlt 5 3
Core-A:1(config)# vlan mlt 2000 3
```

```
Core-A:1(config)# vlan mlt 5 3
Core-A:1(config)# vlan mlt 2000 3
```

3.2.3 CP-Limit

For this configuration step the following CP-Limit parameters will be enabled:

VSP 9000 – Core-A and Core-B		
MLT ID	**CP-Limit Threshold**	**Action**
2 (Edge-SMLT-2)	10000	Shutdown
3 (DC-SMLT-3)	10000	Shutdown

Table 3.2.3 – Core-A / Core-B CP-Limit Parameters

△ Caution – CP-Limit should not be enabled on IST ports.

1 Enable CP-Limit on SMLT 2:

```
Core-A:1(config)# interface mlt 2
Core-A:1(config-if)# cp-limit 10000 shutdown
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface mlt 2
Core-B:1(config-if)# cp-limit 10000 shutdown
Core-B:1(config-if)# exit
```

2 Enable CP-Limit on SMLT 3:

```
Core-A:1(config)# interface mlt 3
Core-A:1(config-if)# cp-limit 10000 shutdown
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface mlt 3
Core-B:1(config-if)# cp-limit 10000 shutdown
Core-B:1(config-if)# exit
```

3.2.4 VLACP

For this configuration step the following VLACP parameters will be enabled:

VSP 9000 – Core-A and Core-B			
VLACP MAC Address	**Port(s)**	**Timeout**	**Timeout Scale**
01:80:c2:00:00:0f	3/1,4/1	Long	3
01:80:c2:00:00:0f	3/3-6, 4/3-6	Short	5

Table 3.2.4 – Core-A / Core-B VLACP Parameters

1 Globally enable VLACP:

```
Core-A:1(config)# vlacp enable
```

```
Core-B:1(config)# vlacp enable
```

2 Enable VLACP on the IST ports:

```
Core-A:1(config)# interface GigabitEthernet 3/1,4/1
Core-A:1(config-if)# vlacp timeout long timeout-scale 3 funcmac-addr 01:80:c2:00:00:0f
Core-A:1(config-if)# vlacp enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface GigabitEthernet 3/1,4/1
Core-B:1(config-if)# vlacp timeout long timeout-scale 3 funcmac-addr 01:80:c2:00:00:0f
Core-B:1(config-if)# vlacp enable
Core-B:1(config-if)# exit
```

3 Enable VLACP on SMLT 2 and SMLT 3 ports:

```
Core-A:1(config)# interface GigabitEthernet 3/3-6,4/3-6
Core-A:1(config-if)# vlacp timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
Core-A:1(config-if)# vlacp enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface GigabitEthernet 3/3-6,4/3-6
Core-B:1(config-if)# vlacp timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
Core-B:1(config-if)# vlacp enable
Core-B:1(config-if)# exit
```

3.2.5 SLPP

For this configuration step the following SLPP parameters will be enabled:

VSP 9000 – Core-A	
VLAN ID	**SLPP Mode**
4	Enabled
5	Enabled
2000	Enabled
Port(s)	**RX Threshold**
3/3-6,4/3-6	5
VSP 9000 – Core-B	
VLAN ID	**SLPP Mode**
4	Enabled
5	Enabled
2000	Enabled
Port(s)	**RX Threshold**
3/3-6,4/3-6	50

Table 3.2.5 – Core-A / Core-B SLPP Parameters

△ Caution – SLPP should not be enabled on IST ports.

1 Globally enable SLPP:

```
Core-A:1(config)# slpp enable
```

```
Core-B:1(config)# slpp enable
```

2 Enable SLPP on SMLT VLANs 4, 5 and 2000:

```
Core-A:1(config)# slpp vid 4
Core-A:1(config)# slpp vid 5
Core-A:1(config)# slpp vid 2000
```

```
Core-B:1(config)# slpp vid 4
Core-B:1(config)# slpp vid 5
Core-B:1(config)# slpp vid 2000
```

2 Enable SLPP on SMLT 2 and SMLT 3 ports:

```
Core-A:1(config)# interface GigabitEthernet 3/3-6,4/3-6
Core-A:1(config-if)# slpp packet-rx
Core-A:1(config-if)# slpp packet-rx-threshold 5
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface GigabitEthernet 3/3-6,4/3-6
Core-B:1(config-if)# slpp packet-rx
Core-B:1(config-if)# slpp packet-rx-threshold 50
Core-B:1(config-if)# exit
```

3.2.6 Discard Untagged Frames

For this configuration step Discard Untagged Frames will be enabled on all IST and SMLT ports:

VSP 9000 – Core-A and Core-B	
Port(s)	**Discard Untagged Frames**
3/1,3/3-6,4/1,4/3-6	Enabled

Table 3.2.6 – Core-A / Core-B Discard Untagged Frame Parameters

1 Enable Discard Untagged Frames on all IST and SMLT ports:

```
Core-A:1(config)# interface GigabitEthernet 3/1,3/3-6,4/1,4/3-6
Core-A:1(config-if)# untagged-frames-discard
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface GigabitEthernet 3/1,3/3-6,4/1,4/3-6
Core-B:1(config-if)# untagged-frames-discard
Core-B:1(config-if)# exit
```

3.2.7 Quality of Service

For this configuration step QoS trust will be enabled on all IST and SMLT ports:

VSP 9000 – Core-A and Core-B	
Port(s)	**DiffServ Trust**
3/1,3/3-6,4/1,4/3-6	Enabled

Table 3.2.7 – Core-A / Core-B QoS Parameters

1 Enable QoS Trusted Interfaces on all IST and SMLT ports:

```
Core-A:1(config)# interface GigabitEthernet 3/1,3/3-6,4/1,4/3-6
Core-A:1(config-if)# enable-diffserv enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface GigabitEthernet 3/1,3/3-6,4/1,4/3-6
Core-B:1(config-if)# enable-diffserv enable
Core-B:1(config-if)# exit
```

3.2.8 RSMLT

For this configuration step the following RSMLT parameters will be enabled:

VSP 9000 – Core-A			
VLAN ID	**IP Address**	**Hold Down Timer**	**Hold Up Timer**
4 (Edge-SMLT-2)	10.0.4.1/24	60	180
5 (DC-SMLT-3)	10.0.5.1/24	60	180
VSP 9000 – Core-B			
VLAN ID	**IP Address**	**Hold Down Timer**	**Hold Up Timer**
4 (Edge-SMLT-2)	10.0.4.2/24	60	180
5 (DC-SMLT-3)	10.0.5.2/24	60	180

Table 3.2.8 – Core-A / Core-B RSMLT Parameters

1	Enable RSMLT on VLANs 4 & 5:

```
Core-A:1(config)# interface vlan 4
Core-A:1(config-if)# ip rsmlt
Core-A:1(config-if)# exit
Core-A:1(config)# interface vlan 5
Core-A:1(config-if)# ip rsmlt
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 4
Core-B:1(config-if)# ip rsmlt
Core-B:1(config-if)# exit
Core-B:1(config)# interface vlan 5
Core-B:1(config-if)# ip rsmlt
Core-B:1(config-if)# exit
```

3.2.9 OSPF

For this configuration step the following OSPF parameters will be enabled:

VSP 9000 – Core-A	
IP Interface	**Priority**
Loopback 1 (1.0.0.1/32)	Default
VLAN 4 (Edge-SMLT-2)	50
VLAN 5 (DC-SMLT-3)	0
VSP 9000 – Core-B	
IP Interface	**Priority**
Loopback 1 (1.0.0.2/32)	Default
VLAN 4 (Edge-SMLT-2)	0
VLAN 5 (DC-SMLT-3)	50

Table 3.2.9 – Core-A / Core-B OSPF Parameters

1 Create loopback interfaces:

```
Core-A:1(config)# interface loopback 1
Core-A:1(config-if)# ip address 1 10.0.0.1/32
Core-A:1(config-if)# ip ospf 1
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface loopback 1
Core-B:1(config-if)# ip address 1 10.0.0.2/32
Core-B:1(config-if)# ip ospf 1
Core-B:1(config-if)# exit
```

2 Globally enable OSPF:

```
Core-A:1(config)# router ospf enable
```

```
Core-B:1(config)# router ospf enable
```

3 Define the OSPF router ID:

```
Core-A:1(config)# router ospf
Core-A:1(config-ospf) router-id 10.0.0.1
Core-A:1(config-ospf)# exit
```

```
Core-B:1(config)# router ospf
Core-B:1(config-ospf) router-id 10.0.0.2
Core-B:1(config-ospf)# exit
```

4 Enable OSPF on VLAN 4 (Edge-SMLT-2):

```
Core-A:1(config)# interface vlan 4
Core-A:1(config-if)# ip ospf enable
Core-A:1(config-if)# ip ospf priority 50
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 4
Core-B:1(config-if)# ip ospf enable
Core-B:1(config-if)# ip ospf priority 0
Core-B:1(config-if)# exit
```

5 Enable OSPF on VLAN 5 (DC-SMLT-3):

```
Core-A:1(config)# interface vlan 5
Core-A:1(config-if)# ip ospf enable
Core-A:1(config-if)# ip ospf priority 0
```

```
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 5
Core-B:1(config-if)# ip ospf enable
Core-B:1(config-if)# ip ospf priority 50
Core-B:1(config-if)# exit
```

3.2.10 PIM-SM

For this configuration step the following PIM-SM parameters will be enabled:

<table>
<tr><th colspan="6">VSP 9000 – Core-A</th></tr>
<tr><th colspan="3">IP Interface</th><th colspan="3">PIM Status</th></tr>
<tr><td colspan="3">VLAN 2 (IST)</td><td colspan="3">Enabled</td></tr>
<tr><td colspan="3">VLAN 4 (Edge-SMLT-2)</td><td colspan="3">Enabled</td></tr>
<tr><td colspan="3">VLAN 5 (DC-SMLT-3)</td><td colspan="3">Enabled</td></tr>
<tr><th colspan="3">IP Interface</th><th colspan="3">BSR Preference</th></tr>
<tr><td colspan="3">Loopback 1</td><td colspan="3">100</td></tr>
<tr><th colspan="2">IP Interface</th><th colspan="2">Multicast Group</th><th colspan="2">RP Candidate</th></tr>
<tr><td colspan="2">Loopback 1</td><td colspan="2">231.0.0.0/8</td><td colspan="2">10.0.0.1</td></tr>
<tr><th colspan="6">VSP 9000 – Core-B</th></tr>
<tr><th colspan="3">IP Interface</th><th colspan="3">PIM Status</th></tr>
<tr><td colspan="3">VLAN 2 (IST)</td><td colspan="3">Enabled</td></tr>
<tr><td colspan="3">VLAN 4 (Edge-SMLT-2)</td><td colspan="3">Enabled</td></tr>
<tr><td colspan="3">VLAN 5 (DC-SMLT-3)</td><td colspan="3">Enabled</td></tr>
<tr><th colspan="3">IP Interface</th><th colspan="3">BSR Preference</th></tr>
<tr><td colspan="3">Loopback 1</td><td colspan="3">50</td></tr>
<tr><th colspan="2">IP Interface</th><th colspan="2">Multicast Group</th><th colspan="2">RP Candidate</th></tr>
<tr><td colspan="2">Loopback 1</td><td colspan="2">231.0.0.0/8</td><td colspan="2">10.0.0.2</td></tr>
</table>

Table 3.2.10 – Core-A / Core-B PIM-SM Parameters

Note – This configuration example assumes the candidate RPs and BSRs are located on the Core-A and Core-B switches.

1 Globally enable PIM:

```
Core-A:1(config)# multicast smlt-square
Core-A:1(config)# ip pim enable
Core-A:1(config)# ip pim fast-joinprune
```

```
Core-B:1(config)# multicast smlt-square
Core-B:1(config)# ip pim enable
Core-B:1(config)# ip pim fast-joinprune
```

2 Enable PIM on the IST VLAN:

```
Core-A:1(config)# interface vlan 2
Core-A:1(config-if)# ip pim enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 2
Core-B:1(config-if)# ip pim enable
Core-B:1(config-if)# exit
```

3 Enable PIM on VLAN 4 (Edge-SMLT-2):

```
Core-A:1(config)# interface vlan 4
Core-A:1(config-if)# ip pim enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 4
Core-B:1(config-if)# ip pim enable
Core-B:1(config-if)# exit
```

4 Enable PIM on VLAN 5 (DC-SMLT-3):

```
Core-A:1(config)# interface vlan 5
Core-A:1(config-if)# ip pim enable
Core-A:1(config-if)# exit
```

```
Core-B:1(config)# interface vlan 5
Core-B:1(config-if)# ip pim enable
Core-B:1(config-if)# exit
```

5 Configure loopback interfaces as the BSR and RP candidates:

```
Core-A:1(config)# interface loopback 1
Core-A:1(config-if)# ip pim 1
Core-A:1(config-if)# ip pim bsr-candidate preference 100
Core-A:1(config-if)# exit
Core-A:1(config)# ip pim rp-candidate group 231.0.0.0 255.0.0.0 rp 10.0.0.1
```

```
Core-B:1(config)# interface loopback 1
Core-B:1(config-if)# ip pim 1
Core-B:1(config-if)# ip pim bsr-candidate preference 50
Core-B:1(config-if)# exit
Core-B:1(config)# ip pim rp-candidate group 231.0.0.0 255.0.0.0 rp 10.0.0.2
```

3.3 Edge Distribution Switch Configuration

This section covers the basic configuration of the Ethernet Routing Switch 8800 edge distribution switch cluster.

Dist-A	Dist-B
VLAN ID:2 (IST) IP: 1.1.1.5/30	VLAN ID:2 (IST) IP: 1.1.1.6/30
VLAN ID: 4 (Edge-SMLT-2) IP: 10.0.4.3/24	VLAN ID: 4 (Edge-SMLT-2) IP: 10.0.4.4/24
VLAN ID: 2101 (User1) IP: 10.21.1.1/24	VLAN ID: 2101 (User1) IP: 10.21.1.2/24
VLAN ID: 2102 (Voice1) IP: 10.21.2.1/24	VLAN ID: 2102 (Voice1) IP: 10.21.2.2/24
VLAN ID: 2201 (User2) IP: 10.22.1.1/24	VLAN ID: 2201 (User2) IP: 10.22.1.2/24
VLAN ID: 2202 (Voice2) IP: 10.22.2.1/24	VLAN ID: 2202 (Voice2) IP: 10.22.2.2/24
VLAN ID: 2000 (Management) IP: 10.20.0.5/24	VLAN ID: 2000 (Management) IP: 10.20.0.6/24
Loopback ID: 1 IP: 1.0.0.5/32	Loopback ID: 1 IP: 1.0.0.6/32
MLT ID: 1 (IST) VLAN IDs: 2,4,2000,2101,2102,2201,2202 Ports: 1/1,2/1	MLT ID: 1 (IST) VLAN IDs: 2,4,2000,2101,2102,2201,2202 Ports: 1/1,2/1
MLT ID: 2 (Edge-SMLT-2) VLAN IDs: 4,2000 Ports: 1/3,1/5,2/3,2/5	MLT ID: 2 (Edge-SMLT-2) VLAN IDs: 4,2000 Ports: 1/3,1/5,2/3,2/5
MLT ID: 10 (SMLT 10) VLAN IDs: 2000,2101,2102 Ports: 1/11	MLT ID: 10 (SMLT 10) VLAN IDs: 2000,2101,2102 Ports: 1/11
SLT ID: 100 VLAN IDs: 2000,2201,2202 Ports: 1/12	SLT ID: 100 VLAN IDs: 2000,2201,2202 Ports: 1/12

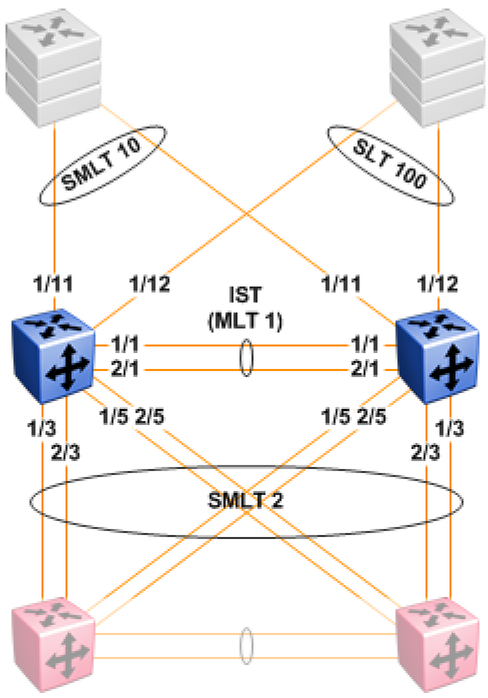

Figure 3.3 – Edge Distribution Switch Cluster Configuration

3.3.1 Virtual LANs

For this configuration step the following VLAN and IP parameters will be enabled:

ERS 8800 – Dist-A		
VLAN Name	**VLAN ID**	**IP Address**
IST	2	1.1.1.5/30
Edge-SMLT-2	4	10.0.4.3/24
Management	2000	10.20.0.5/24
User1	2101	10.21.1.1/24
Voice1	2102	10.21.2.1/24
User2	2201	10.22.1.1/24
Voice2	2202	10.22.2.1/24
ERS 8800 – Dist-B		
VLAN Name	**VLAN ID**	**IP Address**
IST	2	1.1.1.6/30
Edge-SMLT-2	4	10.0.4.4/24
Management	2000	10.20.0.6/24
User1	2101	10.21.1.2/24
Voice1	2102	10.21.2.2/24
User2	2201	10.22.1.2/24
Voice2	2202	10.22.2.2/24

Table 3.3.1 – Dist-A / Dist-B VLAN and IP Interface Parameters

1 Create IST, SMLT and Management VLANs:

```
Dist-A:5# config vlan 2 create byport 1 name "IST"
Dist-A:5# config vlan 4 create byport 1 name "Edge-SMLT-2"
Dist-A:5# config vlan 4 fdb-entry aging-time 21601
Dist-A:5# config vlan 2000 create byport 1 name "Management"
Dist-A:5# config vlan 2000 fdb-entry aging-time 21601
Dist-A:5# config vlan 2101 create byport 1 name "User1"
Dist-A:5# config vlan 2101 fdb-entry aging-time 21601
Dist-A:5# config vlan 2102 create byport 1 name "Voice1"
Dist-A:5# config vlan 2102 fdb-entry aging-time 21601
Dist-A:5# config vlan 2201 create byport 1 name "User2"
Dist-A:5# config vlan 2201 fdb-entry aging-time 21601
Dist-A:5# config vlan 2202 create byport 1 name "Voice2"
Dist-A:5# config vlan 2202 fdb-entry aging-time 21601
```

```
Dist-B:5# config vlan 2 create byport 1 name "IST"
Dist-B:5# config vlan 4 create byport 1 name "Edge-SMLT-2"
Dist-B:5# config vlan 4 fdb-entry aging-time 21601
Dist-B:5# config vlan 2000 create byport 1 name "Management"
Dist-B:5# config vlan 2000 fdb-entry aging-time 21601
Dist-B:5# config vlan 2101 create byport 1 name "User1"
Dist-B:5# config vlan 2101 fdb-entry aging-time 21601
Dist-B:5# config vlan 2102 create byport 1 name "Voice1"
Dist-B:5# config vlan 2102 fdb-entry aging-time 21601
Dist-B:5# config vlan 2201 create byport 1 name "User2"
Dist-B:5# config vlan 2201 fdb-entry aging-time 21601
Dist-B:5# config vlan 2202 create byport 1 name "Voice2"
Dist-B:5# config vlan 2202 fdb-entry aging-time 21601
```

2 Add IP Addresses to the IST and SMLT VLANs:

```
Dist-A:5# config vlan 2 ip create 1.1.1.5/255.255.255.252
Dist-A:5# config vlan 4 ip create 10.0.4.3/255.255.255.0
Dist-A:5# config vlan 2000 ip create 10.20.0.5/255.255.255.0
Dist-A:5# config vlan 2101 ip create 10.21.1.1/255.255.255.0
Dist-A:5# config vlan 2102 ip create 10.21.2.1/255.255.255.0
Dist-A:5# config vlan 2201 ip create 10.22.1.1/255.255.255.0
Dist-A:5# config vlan 2202 ip create 10.22.2.1/255.255.255.0
```

```
Dist-B:5# config vlan 2 ip create 1.1.1.6/255.255.255.252
Dist-B:5# config vlan 4 ip create 10.0.4.4/255.255.255.0
Dist-B:5# config vlan 2000 ip create 10.20.0.6/255.255.255.0
Dist-B:5# config vlan 2101 ip create 10.21.1.2/255.255.255.0
Dist-B:5# config vlan 2102 ip create 10.21.2.2/255.255.255.0
Dist-B:5# config vlan 2201 ip create 10.22.1.2/255.255.255.0
Dist-B:5# config vlan 2202 ip create 10.22.2.2/255.255.255.0
```

3.3.2 Switch Clustering

For this configuration step the following switch clustering parameters will be enabled:

ERS 8800 – Dist-A and Dist-B		
MLT / SLT ID	**Port(s)**	**VLAN Members**
1 (IST)	1/1,2/1	2,4,2000,2101,2102,2201,2202
2 (Edge-SMLT-2)	1/3,1/5,2/3,2/5	4,2000
10 (ERS2500-1)	1/11	2000,2101,2201
100 (ERS4500-1)	1/12	2000,2102,2202

Table 3.3.2 – Dist-A / Dist-B Switch Clustering Parameters

1 Create MLT 1 for IST and assign ports:

```
Dist-A:5# config mlt 1 create
Dist-A:5# config mlt 1 name "IST-Dist-B"
Dist-A:5# config mlt 1 add ports 1/1,2/1
Dist-A:5# config mlt 1 perform-tagging enable
Dist-A:5# config vlan 2 add-mlt 1
```

```
Dist-B:5# config mlt 1 create
Dist-B:5# config mlt 1 name "IST-Dist-A"
Dist-B:5# config mlt 1 add ports 1/1,2/1
Dist-B:5# config mlt 1 perform-tagging enable
Dist-B:5# config vlan 2 add-mlt 1
```

2 Create IST:

```
Dist-A:5# config mlt 1 ist create ip 1.1.1.6 vlan-id 2
Dist-A:5# config mlt 1 ist enable
```

```
Dist-B:5# config mlt 1 ist create ip 1.1.1.5 vlan-id 2
Dist-B:5# config mlt 1 ist enable
```

3 Create SMLT 2 to VSP 9000 Core-A and Core-B Switches:

```
Dist-A:5# config mlt 2 create
Dist-A:5# config mlt 2 add ports 1/3,1/5,2/3,2/5
Dist-A:5# config mlt 2 mcast-distribution enable
Dist-A:5# config mlt 2 name "SMLT-Core-AB"
Dist-A:5# config mlt 2 perform-tagging enable
Dist-A:5# config mlt 2 smlt create smlt-id 2
```

```
Dist-B:5# config mlt 2 create
Dist-B:5# config mlt 2 add ports 1/3,1/5,2/3,2/5
Dist-B:5# config mlt 2 mcast-distribution enable
Dist-B:5# config mlt 2 name "SMLT-Core-AB"
Dist-B:5# config mlt 2 perform-tagging enable
Dist-B:5# config mlt 2 smlt create smlt-id 2
```

4 Create SMLT 10 to ERS 2500 Stack 1:

```
Dist-A:5# config mlt 10 create
Dist-A:5# config mlt 10 add ports 1/11
Dist-A:5# config mlt 10 mcast-distribution enable
Dist-A:5# config mlt 10 name "SMLT-ERS2500-1"
Dist-A:5# config mlt 10 perform-tagging enable
Dist-A:5# config mlt 10 smlt create smlt-id 10
```

```
Dist-B:5# config mlt 10 create
Dist-B:5# config mlt 10 add ports 1/11
Dist-B:5# config mlt 10 mcast-distribution enable
Dist-B:5# config mlt 10 name "SMLT-ERS2500-1"
Dist-B:5# config mlt 10 perform-tagging enable
Dist-B:5# config mlt 10 smlt create smlt-id 10
```

5 Create SLT 100 to ERS 4500 Stack 1:

```
Dist-A:5# config ethernet 1/12 name “SLT-ERS4500-1”
Dist-A:5# config ethernet 1/12 perform-tagging enable
Dist-A:5# config ethernet 1/12 smlt 100 create
```

```
Dist-B:5# config ethernet 1/12 name “SLT-ERS4500-1”
Dist-B:5# config ethernet 1/12 perform-tagging enable
```

```
Dist-B:5# config ethernet 1/12 smlt 100 create
```

6 Add VLANs 4, 2000, 2101, 2102, 2201 and 2202 to the IST:

```
Dist-A:5# config vlan 4 add-mlt 1
Dist-A:5# config vlan 2000 add-mlt 1
Dist-A:5# config vlan 2101 add-mlt 1
Dist-A:5# config vlan 2102 add-mlt 1
Dist-A:5# config vlan 2201 add-mlt 1
Dist-A:5# config vlan 2202 add-mlt 1
```

```
Dist-B:5# config vlan 4 add-mlt 1
Dist-B:5# config vlan 2000 add-mlt 1
Dist-B:5# config vlan 2101 add-mlt 1
Dist-B:5# config vlan 2102 add-mlt 1
Dist-B:5# config vlan 2201 add-mlt 1
Dist-B:5# config vlan 2202 add-mlt 1
```

6 Add VLANs 4 and 2000 to SMLT 2:

```
Dist-A:5# config vlan 4 add-mlt 2
Dist-A:5# config vlan 2000 add-mlt 2
```

```
Dist-B:5# config vlan 4 add-mlt 2
Dist-B:5# config vlan 2000 add-mlt 2
```

7 Add VLANs 2000, 2101, 2102 to SMLT 10:

```
Dist-A:5# config vlan 2000 add-mlt 10
Dist-A:5# config vlan 2101 add-mlt 10
Dist-A:5# config vlan 2102 add-mlt 10
```

```
Dist-B:5# config vlan 2000 add-mlt 10
Dist-B:5# config vlan 2101 add-mlt 10
Dist-B:5# config vlan 2102 add-mlt 10
```

8 Add VLANs 2000, 2201 and 2202 to port 1/12 (SLT 100):

```
Dist-A:5# config vlan 2000 add port 1/12
Dist-A:5# config vlan 2201 add port 1/12
Dist-A:5# config vlan 2202 add port 1/12
```

```
Dist-B:5# config vlan 2000 add port 1/12
Dist-B:5# config vlan 2201 add port 1/12
Dist-B:5# config vlan 2202 add port 1/12
```

3.3.3 Extended CP-Limit

For this configuration step the following Extended CP-Limit parameters will be enabled:

ERS 8800 – Dist-A and Dist-B		
Port(s)	**Threshold Rate**	**Port Action**
1/3,1/5,1/11,1/12,2/3,2/5	50	Softdown

Table 3.3.3 – Dist-A / Dist-B Extended CP-Limit Parameters

△ Caution – Extended CP-Limit should not be enabled on IST ports.

1 Globally enable Extended CP-Limit:

```
Dist-A:5# config sys ext-cp-limit extcplimit enable
```

```
Dist-B:5# config sys ext-cp-limit extcplimit enable
```

2 Enable Extended CP-Limit on SMLT 2, SMLT 10 and SLT 100 ports :

```
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 ext-cp-limit SoftDown threshold-
util-rate 50
```

```
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 ext-cp-limit SoftDown threshold-
util-rate 50
```

3.3.4 VLACP

For this configuration step the following VLACP parameters will be enabled:

ERS 8800 – Dist-A and Dist-B			
VLACP MAC Address	**Port(s)**	**Timeout**	**Timeout Scale**
01:80:c2:00:00:0f	1/1,2/1	Long	3
01:80:c2:00:00:0f	1/3,1/5,1/11-12,2/3,2/5	Short	5

Table 3.3.4 – Dist-A / Dist-B VLACP Parameters

1 Globally enable VLACP:

```
Dist-A:5# config vlacp enable
```

```
Dist-B:5# config vlacp enable
```

2 Enable VLACP on the IST ports:

```
Dist-A:5# config ethernet 1/1,2/1 vlacp macaddress 01:80:c2:00:00:0f
Dist-A:5# config ethernet 1/1,2/1 vlacp slow-periodic-time 10000
Dist-A:5# config ethernet 1/1,2/1 vlacp timeout long
Dist-A:5# config ethernet 1/1,2/1 vlacp timeout-scale 3
Dist-A:5# config ethernet 1/1,2/1 vlacp enable
```

```
Dist-B:5# config ethernet 1/1,2/1 vlacp macaddress 01:80:c2:00:00:0f
Dist-B:5# config ethernet 1/1,2/1 vlacp slow-periodic-time 10000
Dist-B:5# config ethernet 1/1,2/1 vlacp timeout long
Dist-B:5# config ethernet 1/1,2/1 vlacp timeout-scale 3
Dist-B:5# config ethernet 1/1,2/1 vlacp enable
```

3 Enable VLACP on SMLT 2, SMLT 10 and SLT 100 ports:

```
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp macaddress 01:80:c2:00:00:0f
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp fast-periodic-time 500
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp timeout short
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp timeout-scale 5
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp enable
```

```
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp macaddress 01:80:c2:00:00:0f
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp fast-periodic-time 500
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp timeout short
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp timeout-scale 5
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 vlacp enable
```

3.3.5 SLPP

For this configuration step the following SLPP parameters will be enabled:

ERS 8800 – Dist-A	
VLAN ID	**SLPP Mode**
4	Enabled
2000	Enabled
2101	Enabled
2102	Enabled
2201	Enabled
2202	Enabled
Port(s)	**RX Threshold**
1/3,1/5,1/11-12,2/3,2/5	5
ERS 8800 – Dist-B	
VLAN ID	**SLPP Mode**
4	Enabled
2000	Enabled
2101	Enabled
2102	Enabled
2201	Enabled
2202	Enabled
Port(s)	**RX Threshold**
1/3,1/5,1/11-12,2/3,2/5	50

Table 3.3.5 – Dist-A / Dist-B SLPP Parameters

△ Caution – SLPP should not be enabled on IST ports.

1 Enable SLPP on VLANs 4, 2000, 2101, 2102, 2201 and 2202:

```
Dist-A:5# config slpp add 4
Dist-A:5# config slpp add 2000
Dist-A:5# config slpp add 2101
Dist-A:5# config slpp add 2102
Dist-A:5# config slpp add 2201
Dist-A:5# config slpp add 2202
```

```
Dist-B:5# config slpp add 4
Dist-B:5# config slpp add 2000
Dist-B:5# config slpp add 2101
Dist-B:5# config slpp add 2102
Dist-B:5# config slpp add 2201
Dist-B:5# config slpp add 2202
```

2 Enable SLPP on SMLT 2, SMLT 10 and SLT 100 ports:

```
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 slpp packet-rx-threshold 5
Dist-A:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 slpp packet-rx enable
```

```
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 slpp packet-rx-threshold 50
Dist-B:5# config ethernet 1/3,1/5,1/11,1/12,2/3,2/5 slpp packet-rx enable
```

3.3.6 Discard Untagged Frames

For this configuration step Discard Untagged Frames will be enabled on all IST and SMLT ports:

ERS 8800 – Dist-A and Dist-B	
Port(s)	**Discard Untagged Frames**
1/1,1/3,1/5,1/11,1/12,2/1,2/3,2/5	Enabled

Table 3.3.6 – Dist-A / Dist-B Discard Untagged Frame Parameters

1 Enable Discard Untagged Frames on all IST, SMLT and SLT ports:

```
Dist-A:5# config ethernet 1/1,1/3,1/5,1/11,1/12,2/1,2/3,2/5 untagged-frames-discard
enable
```

```
Dist-B:5# config ethernet 1/1,1/3,1/5,1/11,1/12,2/1,2/3,2/5 untagged-frames-discard
enable
```

3.3.7 Quality of Service

For this configuration step QoS trust will be enabled on all IST and SMLT ports:

ERS 8800 – Dist-A and Dist-B	
Port(s)	**DiffServ Trust**
1/1,1/3,1/5,1/11,1/12,2/1,2/3,2/5	Enabled

Table 3.3.7 – Dist-A / Dist-B QoS Parameters

1 Enable QoS Trusted Interfaces on all IST, SMLT and SLT ports:

```
Dist-A:5# config ethernet 1/1,1/3,1/5,1/11,1/12,2/1,2/3,2/5 enable-diffserv true
```

```
Dist-B:5# config ethernet 1/1,1/3,1/5,1/11,1/12,2/1,2/3,2/5 enable-diffserv true
```

3.3.8 RSMLT

For this configuration step the following RSMLT parameters will be enabled:

ERS 8800 – Dist-A			
VLAN ID	**IP Address**	**Hold Down Timer**	**Hold Up Timer**
4 (Edge-SMLT-2)	10.0.4.3/24	60	180
ERS 8800 – Dist-B			
VLAN ID	**IP Address**	**Hold Down Timer**	**Hold Up Timer**
4 (Edge-SMLT-2)	10.0.4.4/24	60	180

Table 3.3.8 – Dist-A / Dist-B RSMLT Parameters

1 Enable QoS Trusted Interfaces on all IST and SMLT ports:

```
Dist-A:5# config vlan 4 ip rsmlt enable
```

```
Dist-B:5# config vlan 4 ip rsmlt enable
```

3.3.9 VRRP Backup / Master

For this configuration step the following VRRP Backup Master parameters will be enabled:

ERS 8800 – Dist-A			
IP Interface	**Virtual Router ID**	**Virtual IP**	**Priority**
10.21.1.1	211	10.21.1.254	200
10.21.2.1	212	10.21.2.254	Default
10.22.1.1	221	10.22.1.254	200
10.22.2.1	222	10.22.2.254	Default
ERS 8800 – Dist-B			
IP Interface	**Virtual Router ID**	**Virtual IP**	**Priority**
10.21.1.2	211	10.21.1.254	Default
10.21.2.2	212	10.21.2.254	200
10.22.1.2	221	10.22.1.254	Default
10.22.2.2	222	10.22.2.254	200

Table 3.3.9 – Dist-A / Dist-B VRRP Parameters

1 Enable VRRP for VLAN 2101:

```
Dist-A:5# config vlan 2101 ip vrrp 211 address 10.21.1.254
Dist-A:5# config vlan 2101 ip vrrp 211 backup-master enable
Dist-A:5# config vlan 2101 ip vrrp 211 priority 200
Dist-A:5# config vlan 2101 ip vrrp 211 adver-int 10
Dist-A:5# config vlan 2101 ip vrrp 211 holddown-timer 60
Dist-A:5# config vlan 2101 ip vrrp 211 enable
```

```
Dist-B:5# config vlan 2101 ip vrrp 211 address 10.21.1.254
Dist-B:5# config vlan 2101 ip vrrp 211 backup-master enable
Dist-B:5# config vlan 2101 ip vrrp 211 adver-int 10
Dist-B:5# config vlan 2101 ip vrrp 211 holddown-timer 60
Dist-B:5# config vlan 2101 ip vrrp 211 enable
```

2 Enable VRRP for VLAN 2102:

```
Dist-A:5# config vlan 2102 ip vrrp 212 address 10.21.2.254
```

```
Dist-A:5# config vlan 2102 ip vrrp 212 backup-master enable
Dist-A:5# config vlan 2102 ip vrrp 212 adver-int 10
Dist-A:5# config vlan 2102 ip vrrp 212 holddown-timer 60
Dist-A:5# config vlan 2102 ip vrrp 212 enable
```

```
Dist-B:5# config vlan 2102 ip vrrp 212 address 10.21.2.254
Dist-B:5# config vlan 2102 ip vrrp 212 backup-master enable
Dist-B:5# config vlan 2102 ip vrrp 212 priority 200
Dist-B:5# config vlan 2102 ip vrrp 212 adver-int 10
Dist-B:5# config vlan 2102 ip vrrp 212 holddown-timer 60
Dist-B:5# config vlan 2102 ip vrrp 212 enable
```

3 Enable VRRP for VLAN 2201:

```
Dist-A:5# config vlan 2201 ip vrrp 221 address 10.22.1.254
Dist-A:5# config vlan 2201 ip vrrp 221 backup-master enable
Dist-A:5# config vlan 2201 ip vrrp 221 priority 200
Dist-A:5# config vlan 2201 ip vrrp 221 adver-int 10
Dist-A:5# config vlan 2201 ip vrrp 221 holddown-timer 60
Dist-A:5# config vlan 2201 ip vrrp 221 enable
```

```
Dist-B:5# config vlan 2201 ip vrrp 221 address 10.22.1.254
Dist-B:5# config vlan 2201 ip vrrp 221 backup-master enable
Dist-B:5# config vlan 2201 ip vrrp 221 adver-int 10
Dist-B:5# config vlan 2201 ip vrrp 221 holddown-timer 60
Dist-B:5# config vlan 2201 ip vrrp 221 enable
```

4 Enable VRRP for VLAN 2202:

```
Dist-A:5# config vlan 2202 ip vrrp 222 address 10.22.2.254
Dist-A:5# config vlan 2202 ip vrrp 222 backup-master enable
Dist-A:5# config vlan 2202 ip vrrp 222 adver-int 10
Dist-A:5# config vlan 2202 ip vrrp 222 holddown-timer 60
Dist-A:5# config vlan 2202 ip vrrp 222 enable
```

```
Dist-B:5# config vlan 2202 ip vrrp 222 address 10.22.2.254
Dist-B:5# config vlan 2202 ip vrrp 222 backup-master enable
Dist-B:5# config vlan 2202 ip vrrp 222 priority 200
Dist-B:5# config vlan 2202 ip vrrp 222 adver-int 10
Dist-B:5# config vlan 2202 ip vrrp 222 holddown-timer 60
Dist-B:5# config vlan 2202 ip vrrp 222 enable
```

3.3.10 OSPF

For this configuration step the following OSPF parameters will be enabled:

ERS 8800 – Dist-A and Dist-B	
IP Interface	**Passive Interface**
VLAN 4 (Edge-SMLT-2)	No
VLAN 2101	Yes
VLAN 2102	Yes
VLAN 2201	Yes
VLAN 2202	Yes

Table 3.3.10 – Dist-A / Dist-B OSPF Parameters

1 Create loopback interfaces:

```
Dist-A:5# config ip circuitless-ip-int 1 create 10.0.0.5/255.255.255.255
Dist-A:5# config ip circuitless-ip-int 1 ospf enable
```

```
Dist-B:5# config ip circuitless-ip-int 1 create 10.0.0.6/255.255.255.255
Dist-B:5# config ip circuitless-ip-int 1 ospf enable
```

2 Define the OSPF router ID:

```
Dist-A:5# config ip ospf router-id 10.0.0.5
```

```
Dist-B:5# config ip ospf router-id 10.0.0.6
```

3 Globally enable OSPF:

```
Dist-A:5# config ip ospf admin-state enable
Dist-A:5# config ip ospf enable
```

```
Dist-B:5# config ip ospf admin-state enable
Dist-B:5# config ip ospf enable
```

4 Enable OSPF on VLAN 4 (Edge-SMLT-2):

```
Dist-A:5# config vlan 4 ip ospf priority 0
Dist-A:5# config vlan 4 ip ospf enable
```

```
Dist-B:5# config vlan 4 ip ospf priority 0
Dist-B:5# config vlan 4 ip ospf enable
```

5 Enable OSPF passive interfaces on VLANs 2101, 2102, 2201 and 2202:

```
Dist-A:5# config vlan 2101 ip ospf interface-type passive
Dist-A:5# config vlan 2101 ip ospf enable
Dist-A:5# config vlan 2102 ip ospf interface-type passive
Dist-A:5# config vlan 2102 ip ospf enable
Dist-A:5# config vlan 2201 ip ospf interface-type passive
Dist-A:5# config vlan 2201 ip ospf enable
Dist-A:5# config vlan 2202 ip ospf interface-type passive
Dist-A:5# config vlan 2202 ip ospf enable
```

```
Dist-B:5# config vlan 2101 ip ospf interface-type passive
Dist-B:5# config vlan 2101 ip ospf enable
Dist-B:5# config vlan 2102 ip ospf interface-type passive
Dist-B:5# config vlan 2102 ip ospf enable
Dist-B:5# config vlan 2201 ip ospf interface-type passive
Dist-B:5# config vlan 2201 ip ospf enable
Dist-B:5# config vlan 2202 ip ospf interface-type passive
Dist-B:5# config vlan 2202 ip ospf enable
```

3.3.11 PIM-SM

For this configuration step the following PIM-SM parameters will be enabled:

ERS 8800 – Dist-A and Dist-B	
IP Interface	**PIM Status**
VLAN 2 (IST)	Enabled
VLAN 4 (Edge-SMLT-2)	Enabled
VLAN 2101	Enabled
VLAN 2102	Enabled
VLAN 2201	Enabled
VLAN 2202	Enabled

Table 3.3.11 – Dist-A / Dist-B PIM-SM Parameters

1 Globally enable PIM:

```
Dist-A:5# sys mcast-smlt square-smlt enable
Dist-A:5# config ip pim enable
Dist-A:5# config ip pim fast-joinprune enable
```

```
Dist-B:5# sys mcast-smlt square-smlt enable
Dist-B:5# config ip pim enable
Dist-B:5# config ip pim fast-joinprune enable
```

2 Enable PIM on the IST VLAN:

```
Dist-A:5# config vlan 2 ip pim enable
```

```
Dist-B:5# config vlan 2 ip pim enable
```

3 Enable PIM on VLAN 4 (Edge-SMLT-2):

```
Dist-A:5# config vlan 4 ip pim enable
```

```
Dist-B:5# config vlan 4 ip pim enable
```

4 Enable PIM on VLANs 2101, 2102, 2201 and 2202:

```
Dist-A:5# config vlan 2101 ip pim enable
Dist-A:5# config vlan 2102 ip pim enable
Dist-A:5# config vlan 2201 ip pim enable
Dist-A:5# config vlan 2202 ip pim enable
```

```
Dist-B:5# config vlan 2101 ip pim enable
Dist-B:5# config vlan 2102 ip pim enable
Dist-B:5# config vlan 2201 ip pim enable
Dist-B:5# config vlan 2202 ip pim enable
```

3.3.12 DHCP Relay

For this configuration step the following DHCP Relay parameters will be enabled:

ERS 8800 – Dist-A		
VLAN	**Interface**	**Forward Path**
2101	10.21.1.1	10.20.10.6
2102	10.21.2.1	10.20.10.6
2201	10.22.1.1	10.20.10.6
2202	10.22.2.1	10.20.10.6
ERS 8800 – Dist-B		
VLAN	**Interface**	**Forward Path**
2101	10.21.1.1	10.20.10.6
2102	10.21.2.1	10.20.10.6
2201	10.22.1.1	10.20.10.6
2202	10.22.2.1	10.20.10.6

Table 3.3.12 – Dist-A / Dist-B DHCP Relay Parameters

Caution – DHCP relay should only be enabled on the real IP interfaces and not the VRRP virtual IP interfaces.

1 Enable DHCP Relay on VLANs 2101, 2102, 2201 and 2202:

```
Dist-A:5# config vlan 2101 ip dhcp-relay enable
Dist-A:5# config vlan 2102 ip dhcp-relay enable
Dist-A:5# config vlan 2201 ip dhcp-relay enable
Dist-A:5# config vlan 2202 ip dhcp-relay enable
```

```
Dist-B:5# config vlan 2101 ip dhcp-relay enable
Dist-B:5# config vlan 2102 ip dhcp-relay enable
Dist-B:5# config vlan 2201 ip dhcp-relay enable
Dist-B:5# config vlan 2202 ip dhcp-relay enable
```

2 Define a DHCP server for VLANs 2101, 2102, 2201 and 2202:

```
Dist-A:5# config ip dhcp-relay create-fwd-path agent 10.21.1.1 server 10.22.2.1 mode
bootp_dhcp state enable

Dist-A:5# config ip dhcp-relay create-fwd-path agent 10.21.2.1 server 10.22.2.1 mode
bootp_dhcp state enable

Dist-A:5# config ip dhcp-relay create-fwd-path agent 10.22.1.1 server 10.22.2.1 mode
bootp_dhcp state enable

Dist-A:5# config ip dhcp-relay create-fwd-path agent 10.22.2.1 server 10.22.2.1 mode
bootp_dhcp state enable
```

```
Dist-B:5# config ip dhcp-relay create-fwd-path agent 10.21.1.2 server 10.22.2.1 mode
bootp_dhcp state enable

Dist-B:5# config ip dhcp-relay create-fwd-path agent 10.21.2.2 server 10.22.2.1 mode
bootp_dhcp state enable

Dist-B:5# config ip dhcp-relay create-fwd-path agent 10.22.1.2 server 10.22.2.1 mode
bootp_dhcp state enable

Dist-B:5# config ip dhcp-relay create-fwd-path agent 10.22.2.2 server 10.22.2.1 mode
bootp_dhcp state enable
```

3.4 Data Center Distribution Switch Configuration

This section covers the basic configuration of the Ethernet Routing Switch 8800 data center distribution switch cluster.

Dist-C
VLAN ID:2 (IST) IP: 1.1.1.9/30
VLAN ID: 5 (DC-SMLT-2) IP: 10.0.4.3/24
VLAN ID: 2000 (Management) IP: 10.20.0.9/24
VLAN ID: 2010 (DataCenter) IP: 10.20.10.1/24
Loopback ID: 1 IP: 1.0.0.9/32
MLT ID: 1 (IST) VLAN IDs: 2,5,2000,2010 Ports: 1/1,2/1
VLAN ID: 5 (DC-SMLT-3) IP: 10.0.5.3/24
MLT ID: 10 (SMLT 10) VLAN IDs: 2000,2010 Ports: 1/11

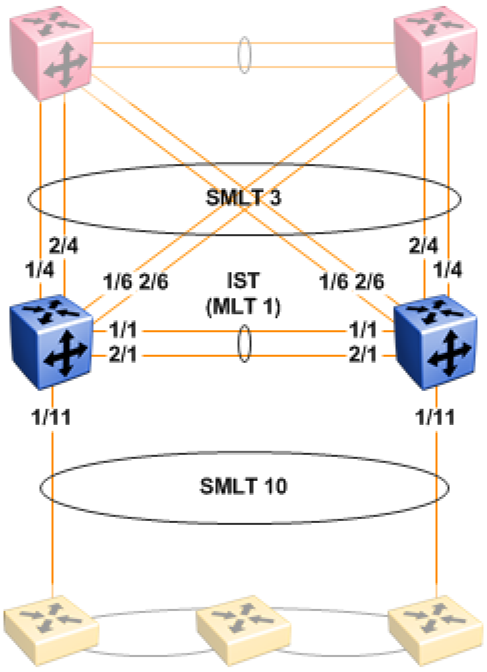

Dist-D
VLAN ID:2 (IST) IP: 1.1.1.10/30
VLAN ID: 5 (DC-SMLT-2) IP: 10.0.4.4/24
VLAN ID: 2000 (Management) IP: 10.20.0.10/24
VLAN ID: 2010 (DataCenter) IP: 10.20.10.2/24
Loopback ID: 1 IP: 1.0.0.10/32
MLT ID: 1 (IST) VLAN IDs: 2,5,2000,2010 Ports: 1/1,2/1
VLAN ID: 5 (DC-SMLT-3) IP: 10.0.5.4/24
MLT ID: 10 (SMLT 10) VLAN IDs: 2000,2010 Ports: 1/11

Figure 3.4 – Data Center Distribution Switch Cluster Configuration

3.4.1 Virtual LANs

For this configuration step the following VLAN and IP parameters will be enabled:

ERS 8800 – Dist-C		
VLAN Name	**VLAN ID**	**IP Address**
IST	2	1.1.1.9/30
DC-SMLT-3	5	10.0.5.3/24
Management	2000	10.20.0.9/24
DataCenter	2010	10.20.10.1/24
ERS 8800 – Dist-D		
VLAN Name	**VLAN ID**	**IP Address**
IST	2	1.1.1.10/30
DC-SMLT-3	5	10.0.5.4/24
Management	2000	10.20.0.10/24
DataCenter	2010	10.20.10.2/24

Table 3.4.1 – Dist-C / Dist-C VLAN and IP Interface Parameters

1 Create IST, SMLT and Management VLANs:

```
Dist-C:5# config vlan 2 create byport 1 name "IST"
Dist-C:5# config vlan 5 create byport 1 name "DC-SMLT-3"
Dist-C:5# config vlan 5 fdb-entry aging-time 21601
Dist-C:5# config vlan 2000 create byport 1 name "Management"
Dist-C:5# config vlan 2000 fdb-entry aging-time 21601
Dist-C:5# config vlan 2010 create byport 1 name "DataCenter"
Dist-C:5# config vlan 2010 fdb-entry aging-time 21601
```

```
Dist-D:5# config vlan 2 create byport 1 name "IST"
Dist-D:5# config vlan 5 create byport 1 name "DC-SMLT-3"
Dist-D:5# config vlan 5 fdb-entry aging-time 21601
Dist-D:5# config vlan 2000 create byport 1 name "Management"
Dist-D:5# config vlan 2000 fdb-entry aging-time 21601
Dist-D:5# config vlan 2010 create byport 1 name "DataCenter"
Dist-D:5# config vlan 2010 fdb-entry aging-time 21601
```

2 Add IP Addresses to the IST and SMLT VLANs:

```
Dist-C:5# config vlan 2 ip create 1.1.1.9/255.255.255.252
Dist-C:5# config vlan 5 ip create 10.0.5.3/255.255.255.0
Dist-C:5# config vlan 2000 ip create 10.20.0.9/255.255.255.0
Dist-C:5# config vlan 2010 ip create 10.20.10.1/255.255.255.0
```

```
Dist-D:5# config vlan 2 ip create 1.1.1.10/255.255.255.252
Dist-D:5# config vlan 5 ip create 10.0.5.4/255.255.255.0
Dist-D:5# config vlan 2000 ip create 10.20.0.10/255.255.255.0
Dist-D:5# config vlan 2010 ip create 10.20.10.2/255.255.255.0
```

3.4.2 Switch Clustering

For this configuration step the following switch clustering parameters will be enabled:

ERS 8800 – Dist-C and Dist-D

MLT ID	Port(s)	VLAN Members
1 (IST)	1/1,2/1	2,4,2000,2010
3 (DC-SMLT-3)	1/4,1/6,2/4,2/6	5,2000
10 (ERS5000-1)	1/11	2000,2010

Table 3.4.2 – Dist-C / Dist-D Switch Clustering Parameters

1 Create MLT 1 for IST and assign ports:

```
Dist-C:5# config mlt 1 create
Dist-C:5# config mlt 1 name "IST-Dist-C"
Dist-C:5# config mlt 1 add ports 1/1,2/1
Dist-C:5# config mlt 1 perform-tagging enable
Dist-C:5# config vlan 2 add-mlt 1
```

```
Dist-D:5# config mlt 1 create
Dist-D:5# config mlt 1 name "IST-Dist-D"
Dist-D:5# config mlt 1 add ports 1/1,2/1
Dist-D:5# config mlt 1 perform-tagging enable
Dist-D:5# config vlan 2 add-mlt 1
```

2 Create IST:

```
Dist-C:5# config mlt 1 ist create ip 1.1.1.10 vlan-id 2
```

```
Dist-C:5# config mlt 1 ist enable
```

```
Dist-D:5# config mlt 1 ist create ip 1.1.1.9 vlan-id 2
Dist-D:5# config mlt 1 ist enable
```

3 Create SMLT 3 to VSP 9000 Core-A and Core-B Switches:

```
Dist-C:5# config mlt 3 create
Dist-C:5# config mlt 3 add ports 1/4,1/6,2/4,2/6
Dist-C:5# config mlt 3 mcast-distribution enable
Dist-C:5# config mlt 3 name "SMLT-Core-AB"
Dist-C:5# config mlt 3 perform-tagging enable
Dist-C:5# config mlt 3 smlt create smlt-id 3
```

```
Dist-D:5# config mlt 3 create
Dist-D:5# config mlt 3 add ports 1/4,1/6,2/4,2/6
Dist-D:5# config mlt 3 mcast-distribution enable
Dist-D:5# config mlt 3 name "SMLT-Core-AB"
Dist-D:5# config mlt 3 perform-tagging enable
Dist-D:5# config mlt 3 smlt create smlt-id 3
```

4 Create SMLT 10 to ERS 5000 Stack 1:

```
Dist-C:5# config mlt 10 create
Dist-C:5# config mlt 10 add ports 1/11
Dist-C:5# config mlt 10 mcast-distribution enable
Dist-C:5# config mlt 10 name "SMLT-ERS5000-1"
Dist-C:5# config mlt 10 perform-tagging enable
Dist-C:5# config mlt 10 smlt create smlt-id 10
```

```
Dist-D:5# config mlt 10 create
Dist-D:5# config mlt 10 add ports 1/11
Dist-D:5# config mlt 10 mcast-distribution enable
Dist-D:5# config mlt 10 name "SMLT-ERS5000-1"
Dist-D:5# config mlt 10 perform-tagging enable
Dist-D:5# config mlt 10 smlt create smlt-id 10
```

5 Add VLANs 5 and 2000 to SMLT 3:

```
Dist-C:5# config vlan 5 add-mlt 3
Dist-C:5# config vlan 2000 add-mlt 3
```

```
Dist-D:5# config vlan 5 add-mlt 3
Dist-D:5# config vlan 2000 add-mlt 3
```

6 Add VLANs 2000 and 2010 to SMLT 10:

```
Dist-C:5# config vlan 2000 add-mlt 10
Dist-C:5# config vlan 2010 add-mlt 10
```

```
Dist-D:5# config vlan 2000 add-mlt 10
Dist-D:5# config vlan 2010 add-mlt 10
```

3.4.3 Extended CP-Limit

For this configuration step the following Extended CP-Limit parameters will be enabled:

ERS 8800 – Dist-C and Dist-D		
Port(s)	**Threshold Rate**	**Port Action**
1/4,1/6,1/11,2/4,2/6	50	Softdown

Table 3.4.3 – Dist-C / Dist-D Extended CP-Limit Parameters

△ Caution – Extended CP-Limit should not be enabled on IST ports.

1 Globally enable Extended CP-Limit:

```
Dist-C:5# config sys ext-cp-limit extcplimit enable
```

```
Dist-D:5# config sys ext-cp-limit extcplimit enable
```

2 Enable Extended CP-Limit on SMLT 2, SMLT 10 and SLT 100 ports :

```
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 ext-cp-limit SoftDown threshold-util-rate 50
```

```
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 ext-cp-limit SoftDown threshold-util-rate 50
```

3.4.4 VLACP

For this configuration step the following VLACP parameters will be enabled:

ERS 8800 – Dist-C and Dist-D			
VLACP MAC Address	**Port(s)**	**Timeout**	**Timeout Scale**
01:80:c2:00:00:0f	1/1,2/1	Long	3
01:80:c2:00:00:0f	1/4,1/6,1/11,2/4,2/6	Short	5

Table 3.4.4 – Dist-C / Dist-D VLACP Parameters

1 Globally enable VLACP:

```
Dist-C:5# config vlacp enable
```

```
Dist-D:5# config vlacp enable
```

2 Enable VLACP on the IST ports:

```
Dist-C:5# config ethernet 1/1,2/1 vlacp macaddress 01:80:c2:00:00:0f
Dist-C:5# config ethernet 1/1,2/1 vlacp slow-periodic-time 10000
Dist-C:5# config ethernet 1/1,2/1 vlacp timeout long
Dist-C:5# config ethernet 1/1,2/1 vlacp timeout-scale 3
Dist-C:5# config ethernet 1/1,2/1 vlacp enable
```

```
Dist-D:5# config ethernet 1/1,2/1 vlacp macaddress 01:80:c2:00:00:0f
Dist-D:5# config ethernet 1/1,2/1 vlacp slow-periodic-time 10000
Dist-D:5# config ethernet 1/1,2/1 vlacp timeout long
Dist-D:5# config ethernet 1/1,2/1 vlacp timeout-scale 3
Dist-D:5# config ethernet 1/1,2/1 vlacp enable
```

3 Enable VLACP on SMLT 3 and SMLT 10 ports:

```
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp macaddress 01:80:c2:00:00:0f
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp fast-periodic-time 500
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp timeout short
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp timeout-scale 5
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp enable
```

```
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp macaddress 01:80:c2:00:00:0f
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp fast-periodic-time 500
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp timeout short
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp timeout-scale 5
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 vlacp enable
```

3.4.5 SLPP

For this configuration step the following SLPP parameters will be enabled:

ERS 8800 – Dist-C	
VLAN ID	**SLPP Mode**
4	Enabled
2000	Enabled
2010	Enabled
Port(s)	**RX Threshold**
1/4,1/6,1/11,2/4,2/6	5
ERS 8800 – Dist-D	
VLAN ID	**SLPP Mode**
4	Enabled
2000	Enabled
2010	Enabled
Port(s)	**RX Threshold**
1/4,1/6,1/11,2/4,2/6	50

Table 3.4.5 – Dist-C / Dist-D SLPP Parameters

△ Caution – SLPP should not be enabled on IST ports.

1 Enable SLPP on VLANs 5 and 2010:

```
Dist-C:5# config slpp add 5
Dist-C:5# config slpp add 2000
Dist-C:5# config slpp add 2010
```

```
Dist-D:5# config slpp add 5
Dist-D:5# config slpp add 2000
Dist-D:5# config slpp add 2010
```

2 Enable SLPP on SMLT 3 and SMLT 10 ports:

```
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 slpp packet-rx-threshold 5
Dist-C:5# config ethernet 1/4,1/6,1/11,2/4,2/6 slpp packet-rx enable
```

```
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 slpp packet-rx-threshold 50
Dist-D:5# config ethernet 1/4,1/6,1/11,2/4,2/6 slpp packet-rx enable
```

3.4.6 Discard Untagged Frames

For this configuration step Discard Untagged Frames will be enabled on all IST and SMLT ports:

ERS 8800 – Dist-C and Dist-D	
Port(s)	**Discard Untagged Frames**
1/1,1/4,1/6,1/11,2/1,2/4,2/6	Enabled

Table 3.4.6 – Dist-A / Dist-B Discard Untagged Frame Parameters

1 Enable Discard Untagged Frames on all IST and SMLT ports:

```
Dist-C:5# config Ethernet 1/1,1/4,1/6,1/11,2/1,2/4,2/6 untagged-frames-discard enable
```

```
Dist-D:5# config Ethernet 1/1,1/4,1/6,1/11,2/1,2/4,2/6 untagged-frames-discard enable
```

3.4.7 Quality of Service

For this configuration step QoS trust will be enabled on all IST and SMLT ports:

ERS 8800 – Dist-C and Dist-D	
Port(s)	**DiffServ Trust**
1/1,1/4,1/6,1/11,2/1,2/4,2/6	Enabled

Table 3.4.7 – Dist-A / Dist-B QoS Parameters

1 Enable QoS Trusted Interfaces on all IST and SMLT ports:

```
Dist-C:5# config ethernet 1/1,1/4,1/6,1/11,2/1,2/4,2/6 enable-diffserv true
```

```
Dist-D:5# config ethernet 1/1,1/4,1/6,1/11,2/1,2/4,2/6 enable-diffserv true
```

3.4.8 RSMLT and RSMLT Layer 2 Edge

For this configuration step the following RSMLT parameters will be enabled:

ERS 8800 – Dist-A			
VLAN ID	**IP Address**	**Hold Down Timer**	**Hold Up Timer**
5 (DC-SMLT-3)	10.0.4.3/24	60	180
2010	10.20.10.1/24	60	9999 (Infinity)
ERS 8800 – Dist-B			
VLAN ID	**IP Address**	**Hold Down Timer**	**Hold Up Timer**
5 (DC-SMLT-3)	10.0.4.4/24	60	180
2010	10.20.10.2/24	60	9999 (Infinity)

Table 3.4.8 – Dist-A / Dist-B RSMLT Parameters

1 Enable QoS Trusted Interfaces on all IST and SMLT ports:

```
Dist-C:5# config vlan 5 ip rsmlt enable
Dist-C:5# config vlan 2010 ip rsmlt enable
Dist-C:5# config vlan 2010 ip rsmlt holdup-timer 9999
Dist-C:5# config ip rsmlt rsmlt-edge-support enable
```

```
Dist-D:5# config vlan 5 ip rsmlt enable
Dist-D:5# config vlan 2010 ip rsmlt enable
Dist-D:5# config vlan 2010 ip rsmlt holdup-timer 9999
Dist-D:5# config ip rsmlt rsmlt-edge-support enable
```

3.4.9 OSPF

For this configuration step the following OSPF parameters will be enabled:

ERS 8800 – Dist-C and Dist-D	
IP Interface	**Passive Interface**
VLAN 5 (DC-SMLT-3)	No
VLAN 2010	Yes

Table 3.4.9 – Dist-C / Dist-D OSPF Parameters

1 Create loopback interfaces:

```
Dist-C:5# config ip circuitless-ip-int 1 create 10.0.0.9/255.255.255.255
Dist-C:5# config ip circuitless-ip-int 1 ospf enable
```

```
Dist-D:5# config ip circuitless-ip-int 1 create 10.0.0.10/255.255.255.255
Dist-D:5# config ip circuitless-ip-int 1 ospf enable
```

2 Define the OSPF router ID:

```
Dist-C:5# config ip ospf router-id 10.0.0.9
```

```
Dist-D:5# config ip ospf router-id 10.0.0.10
```

3 Globally enable OSPF:

```
Dist-C:5# config ip ospf admin-state enable
Dist-C:5# config ip ospf enable
```

```
Dist-D:5# config ip ospf admin-state enable
Dist-D:5# config ip ospf enable
```

4 Enable OSPF on VLAN 5 (DC-SMLT-3):

```
Dist-C:5# config vlan 5 ip ospf priority 0
Dist-C:5# config vlan 5 ip ospf enable
```

```
Dist-D:5# config vlan 5 ip ospf priority 0
Dist-D:5# config vlan 5 ip ospf enable
```

5 Enable OSPF passive interface on VLAN 2010:

```
Dist-C:5# config vlan 2010 ip ospf interface-type passive
Dist-C:5# config vlan 2010 ip ospf enable
```

```
Dist-D:5# config vlan 2010 ip ospf interface-type passive
Dist-D:5# config vlan 2010 ip ospf enable
```

3.4.10 PIM-SM

For this configuration step the following PIM-SM parameters will be enabled:

ERS 8800 – Dist-C and Dist-D	
IP Interface	**PIM Status**
VLAN 2 (IST)	Enabled
VLAN 5 (DC-SMLT-3)	Enabled
VLAN 2010	Enabled

Table 3.4.10 – Dist-C / Dist-D PIM-SM Parameters

1 Globally enable PIM:

```
Dist-C:5# sys mcast-smlt square-smlt enable
Dist-C:5# config ip pim enable
Dist-C:5# config ip pim fast-joinprune enable
```

```
Dist-D:5# sys mcast-smlt square-smlt enable
Dist-D:5# config ip pim enable
Dist-D:5# config ip pim fast-joinprune enable
```

2 Enable PIM on the IST VLAN:

```
Dist-C:5# config vlan 2 ip pim enable
```

```
Dist-D:5# config vlan 2 ip pim enable
```

3 Enable PIM on VLAN 5 (DC-SMLT-3):

```
Dist-C:5# config vlan 5 ip pim enable
```

```
Dist-D:5# config vlan 5 ip pim enable
```

4 Enable PIM on VLANs 2010:

```
Dist-C:5# config vlan 2010 ip pim enable
```

```
Dist-D:5# config vlan 2010 ip pim enable
```

3.5 ERS 2500 / 4500 Edge Switch Configuration

This section covers the basic configuration of the Ethernet Routing Switch 2500 / 4500 edge switches.

ERS 2500

VLAN ID: 2000 (Management)

IP: 10.20.0.13/24

MLT ID: 1 (SMLT 10)

VLAN ID: 2101,2102,2000

Ports: 1/50,3/50

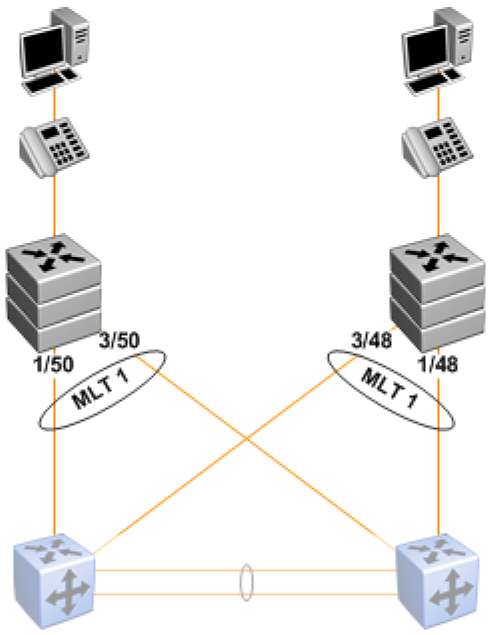

ERS 4500

VLAN ID: 2000 (Management)

IP: 10.20.0.14/24

MLT ID: 1 (SLT 100)

VLAN ID: 2201,2202,2000

Ports: 1/48,3/48

Figure 3.5 – ERS 2500 / 4500 Edge Switch Configuration

3.5.1 Virtual LANs

For this configuration step the following VLAN and IP parameters will be enabled:

ERS 2500	
VLAN Name	**VLAN ID**
Management	2000
User1	2101
Voice1	2102
ERS 4500	
VLAN Name	**VLAN ID**
Management	2000
User2	2201
Voice2	2202

Table 3.5.1 – ERS 2500 / 4500 VLAN Parameters

Note – This configuration assumes ADAC with LLDP/LLDP-MED is used to automatically provision the Voice VLAN when an Avaya IP Phone is connected to the edge switch.

1 Create VLANs:

```
ERS2500-1(config)# vlan create 2000 name Management type port
ERS2500-1(config)# vlan create 2101 name User1 type port
ERS2500-1(config)# vlan create 2102 name Voice1 type port
```

```
ERS4500-1(config)# vlan create 2000 name Management type port
ERS4500-1(config)# vlan create 2201 name User2 type port
ERS4500-1(config)# vlan create 2202 name Voice2 type port
```

2 Configure VLAN Membership and Port Tagging:

```
ERS2500-1(config)# vlan ports 1/50,3/50 tagging tagall filter-untagged-frame enable
ERS2500-1(config)# vlan members remove 1 all
ERS2500-1(config)# vlan members add 2000 1/50,3/50
ERS2500-1(config)# vlan members add 2101 1/1-3/50
```

```
ERS4500-1(config)# vlan ports 1/48,3/48 tagging tagall filter-untagged-frame enable
ERS4500-1(config)# vlan members remove 1 all
ERS4500-1(config)# vlan members add 2000 1/48,3/48
ERS4500-1(config)# vlan members add 2201 1/1-3/48
```

Tip – The ERS 2500 and ERS 4500 switches use the VLAN configuration mode of strict (default setting). In this mode, you must first remove port members from the default VLAN 1 prior to adding these port members to a new VLAN. The VLAN configuration mode is set by using the command vlan configcontrol <automatic|autopvid|flexible|strict>.

3.5.2 Multilink Trunking (MLT)

For this configuration step the following MLT parameters will be enabled:

ERS 2500		
MLT ID	**Port(s)**	**Learning**
1	1/50,3/50	Disabled
ERS 4500		
MLT ID	**Port(s)**	**Learning**
1	1/48,3/48	Disabled

Table 3.5.2 – ERS 2500 / 4500 MLT Parameters

1 Create MLT:

```
ERS2500-1(config)# mlt 1 name Dist-AB member 1/50,3/50 learning disabled
ERS2500-1(config)# mlt 1 loadbalance advance
ERS2500-1(config)# mlt 1 enable

ERS4500-1(config)# mlt 1 name Dist-AB member 1/48,3/48 learning disabled
ERS2500-1(config)# mlt 1 loadbalance advance
ERS4500-1(config)# mlt 1 enable
```

Tip – When configuring the MLT group, Spanning Tree must be disabled on the MLT. The *learning disable* on the MLT command disables Spanning Tree for the MLT.

Note – In an RSMLT square/full mesh core topology, Avaya recommends to set the MLT algorithm on Avaya edge stackable switches to advance for IP based traffic. Tests have shown that this will increase performance over the default MLT setting of basic.

3.5.3 Management

For this configuration step the following Management parameters will be enabled:

ERS 2500		
VLAN Name	**VLAN ID**	**IP Address**
Management	2000	10.20.0.13/24
ERS 4500		
VLAN Name	**VLAN ID**	**IP Address**
Management	2000	10.20.0.14/24

Table 3.5.3 – ERS 2500 / 4500 Management Parameters

1 Specify a Management VLAN:

```
ERS2500-1(config)# vlan mgmt 2000
```

```
ERS4500-1(config)# vlan mgmt 2000
```

2 Create a Management IP Address:

```
ERS2500-1(config)# interface vlan 2000
ERS2500-1(config-if)# ip address 10.20.0.13 255.255.255.0
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface vlan 2000
ERS4500-1(config-if)# ip address 10.20.0.14 255.255.255.0
ERS4500-1(config-if)# exit
```

3.5.4 VLACP

For this configuration step the following VLACP parameters will be enabled:

ERS 2500				
VLACP MAC Address	**Port(s)**	**Timeout**	**Fast Periodic Time**	**Timeout Scale**
01:80:c2:00:00:0f	1/50,3/50	Short	500	5
ERS 4500				
VLACP MAC Address	**Port(s)**	**Timeout**	**Fast Periodic Time**	**Timeout Scale**
01:80:c2:00:00:0f	1/48,3/48	Short	500	5

Table 3.5.4 – ERS 2500 / 4500 VLACP Parameters

1 Globally enable VLACP:

```
ERS2500-1(config)# vlacp macaddress 01:80:c2:00:00:0f
ERS2500-1(config)# enable
```

```
ERS4500-1(config)# vlacp macaddress 01:80:c2:00:00:0f
ERS4500-1(config)# enable
```

2 Enable VLACP on the MLT 1 ports:

```
ERS2500-1(config)# interface fastEthernet all
ERS2500-1(config-if)# vlacp port 1/50,3/50 timeout short
ERS2500-1(config-if)# vlacp port 1/50,3/50 fast-periodic-time 500
ERS2500-1(config-if)# vlacp port 1/50,3/50 timeout-scale 5
ERS2500-1(config-if)# vlacp port 1/50,3/50 enable
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet all
ERS4500-1(config-if)# vlacp port 1/50,3/50 timeout short
ERS4500-1(config-if)# vlacp port 1/50,3/50 fast-periodic-time 500
ERS4500-1(config-if)# vlacp port 1/50,3/50 timeout-scale 5
ERS4500-1(config-if)# vlacp port 1/50,3/50 enable
ERS4500-1(config-if)# exit
```

3.5.5 Spanning Tree Protocol / BPDU Filtering

For this configuration step the following Spanning Tree Protocol parameters will be enabled:

ERS 2500			
Port(s)	**Learning Mode**	**BPDU Filtering Timeout**	**BPDU Filtering Enable**
1/1-49,2/1-50,3/1-49	Fast	0	Enable
ERS 4500			
Port(s)	**Learning Mode**	**BPDU Filtering Timeout**	**BPDU Filtering Enable**
1/1-47,2/1-48,3/1-47	Fast	0	Enable

Table 3.5.5 – ERS 2500 / 4500 Spaning Tree Protocol Parameters

△ Caution – Spanning Tree Protocol and BPDU filtering should only be enabled on edge ports and not MLT ports.

1 Enable Fast Start Spanning Tree Protocol and BPDU Filtering on Edge Ports:

```
ERS2500-1(config)# interface fastEthernet 1/1-49,2/1-50,3/1-49
ERS2500-1(config-if)# spanning-tree learning fast
ERS2500-1(config-if)# spanning-tree bpdu-filtering timeout 0
ERS2500-1(config-if)# spanning-tree bpdu-filtering enable
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet 1/1-47,2/1-48,3/1-47
ERS4500-1(config-if)# spanning-tree learning fast
ERS4500-1(config-if)# spanning-tree bpdu-filtering timeout 0
ERS4500-1(config-if)# spanning-tree bpdu-filtering enable
ERS4500-1(config-if)# exit
```

3.5.6 Rate Limiting

For this configuration step the following Rale Limiting parameters will be enabled:

ERS 2500	
Port(s)	**Broadcast / Multicast Rate (Packets / Second)**
1/1-1/49,2/1-3/50,3/1-3/49	262,143
ERS 4500	
Port(s)	**Broadcast / Multicast Rate (Percent)**
1/1-47,2/1-48,3/1-47	10

Table 3.5.4 – ERS 2500 / 4500 Spaning Tree Protocol Parameters

1 Enable Broadcast and Multicast Rate Limiting on Edge ports:

```
ERS2500-1(config)# interface fastEthernet all
ERS2500-1(config-if)# rate-limit port 1/1-49,2/1-50,3/1-49 both 262143
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet all
ERS4500-1(config-if)# rate-limit port 1/1-47,2/1-48,3/1-47 both 10
ERS4500-1(config-if)# exit
```

3.5.7 ADAC

For this configuration step the following ADAC parameters will be enabled:

ERS 2500			
Port(s)	**Op Mode**	**Voice VLAN**	**Uplink Port**
1/1-1/49,2/1-3/50,3/1-3/49	tagged-frames	2102	1/50
ERS 4500			
Port(s)	**Op Mode**	**Voice VLAN**	**Uplink Port**
1/1-47,2/1-48,3/1-47	tagged-frames	2202	1/48

Table 3.5.7 – ERS 2500 / 4500 ADAC Parameters

Tip – Only one port in the MLT trunk needs to be defined as an ADAC uplink port. The second port in the trunk will be added automatically.

Note – The Data VLAN must be created and provisioned on the port BEFORE enabling ADAC. If the Data VLAN is added to the port after enabling ADAC, once ADAC is disabled, the PVID of the port will be reset to 1 (default VLAN).

1 Configure the ADAC op-mode, define a Voice VLAN and Uplink port:

```
ERS2500-1(config)# adac enable op-mode tagged-frames voice-vlan 2102 uplink-port 1/50
```

```
ERS4500-1(config)# adac enable op-mode tagged-frames voice-vlan 2202 uplink-port 1/48
```

2 Enable ADAC on Edge ports:

```
ERS2500-1(config)# interface fastEthernet all
ERS2500-1(config-if)# adac port 1/1-49,2/1-50,3/1-49 enable
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet all
ERS4500-1(config-if)# adac port 1/1-47,2/1-48,3/1-47 enable
ERS2500-1(config-if)# exit
```

3 Configure ADAC operation mode and enable ADAC:

```
ERS2500-1(config-if)# adac detection lldp
ERS2500-1(config-if)# no adac delection mac
ERS2500-1(config-if)# adac enable
```

```
ERS2500-1(config-if)# adac detection lldp
ERS2500-1(config-if)# no adac delection mac
ERS4500-1(config-if)# adac enable
```

4 Enable LLDP-MED:

```
ERS2500-1(config-if)# lldp tx-tlv local-mgmt-addr port-desc sys-cap sys-desc sys-name
ERS2500-1(config-if)# lldp status txAndRx config-notification
ERS2500-1(config-if)# lldp tx-tlv med extendedPSE med-capabilities network-policy
ERS2500-1(config-if)# exit
```

```
ERS2500-1(config-if)# lldp tx-tlv local-mgmt-addr port-desc sys-cap sys-desc sys-name
ERS2500-1(config-if)# lldp status txAndRx config-notification
ERS2500-1(config-if)# lldp tx-tlv med extendedPSE med-capabilities network-policy
ERS2500-1(config-if)# exit
```

3.5.8 Enable IP Source Guard and DAI

For this configuration step the following ADAC parameters will be enabled:

ERS 2500			
Port(s)	**DHCP Snooping**	**DAI**	**IP Source Guard**
1/50,3/50	Trusted	Trusted	Disabled
1/1-1/49,2/1-3/50,3/1-3/49	Untrusted	Untrusted	Enabled
ERS 4500			
Port(s)	**DHCP Snooping**	**DAI**	**IP Source Guard**
1/48,3/48	Trusted	Trusted	Disabled
1/1-47,2/1-48,3/1-47	Untrusted	Untrusted	Enabled

Table 3.5.8 – ERS 2500 / 4500 DHCP Smooping and DAI Parameters

△ Caution – IP Source Guard and DAI should not be enabled on MLT ports.

1 Globally enable DHCP Snooping:

```
ERS2500-1(config)# ip dhcp-snooping enable
```

```
ERS4500-1(config)# ip dhcp-snooping enable
```

2 Enable DHCP Snooping on User and Voice VLANs:

```
ERS2500-1(config)# ip dhcp-snooping vlan 2101
ERS2500-1(config)# ip dhcp-snooping vlan 2102
```

```
ERS4500-1(config)# ip dhcp-snooping vlan 2201
ERS4500-1(config)# ip dhcp-snooping vlan 2202
```

3 Enable DHCP trust on MLT 1 ports:

```
ERS2500-1(config)# interface fastEthernet 1/50,3/50
ERS2500-1(config-if)# ip dhcp-snooping trusted
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet 1/48,3/48
ERS4500-1(config-if)# ip dhcp-snooping trusted
ERS4500-1(config-if)# exit
```

4 Enable Dynamic ARP Inspection on User and Voice VLANs:

```
ERS2500-1(config)# ip arp-inspection vlan 2101
ERS2500-1(config)# ip arp-inspection vlan 2102
```

```
ERS4500-1(config)# ip arp-inspection vlan 2201
ERS4500-1(config)# ip arp-inspection vlan 2202
```

5 Enable DAI trust on MLT 1 ports:

```
ERS2500-1(config)# interface fastEthernet 1/50,3/50
ERS2500-1(config-if)# ip arp-inspection trusted
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet 1/50,3/50
ERS4500-1(config-if)# ip arp-inspection trusted
ERS4500-1(config-if)# exit
```

6 Enable IP Source Guard on Edge ports:

```
ERS2500-1(config)# interface fastEthernet all
ERS2500-1(config-if)# ip verify source interface fastEthernet 1/1-1/49,2/1-3/50,3/1-3/49
ERS2500-1(config-if)# exit
```

```
ERS4500-1(config)# interface fastEthernet all
ERS4500-1(config-if)# ip verify source interface fastEthernet 1/1-1/47,2/1-3/48,3/1-3/47
ERS4500-1(config-if)# exit
```

3.5.9 IGMP Snooping and Proxy

For this configuration step the following IGMP parameters will be enabled:

ERS 2500		
VLAN ID	**IGMP Snooping**	**IGMP Proxy**
2101	Yes	Yes
ERS 4500		
VLAN ID	**IGMP Snooping**	**IGMP Proxy**
2201	Yes	Yes

Table 3.5.9 – ERS 2500 / 4500 IGMP Parameters

Caution – IGMP should only be enabled on the User VLANs and not the Voice VLANs.

1 Enable IGMP Snooping and Proxy on User VLANs:

```
ERS2500-1(config)# vlan igmp 2101 snooping enable
ERS2500-1(config)# vlan igmp 2101 proxy enable
```

```
ERS4500-1(config)# vlan igmp 2201 snooping enable
ERS4500-1(config)# vlan igmp 2201 proxy enable
```

3.6 ERS 5000 Data Center Switch Configuration

This section covers the basic configuration of the Ethernet Routing Switch 5000 data center switches.

ERS 5000
VLAN ID: 2000 (Management) IP: 10.20.0.17/24
MLT ID: 1 (SMLT 10) VLAN ID: 2000,2010 Ports: 1/50,3/50

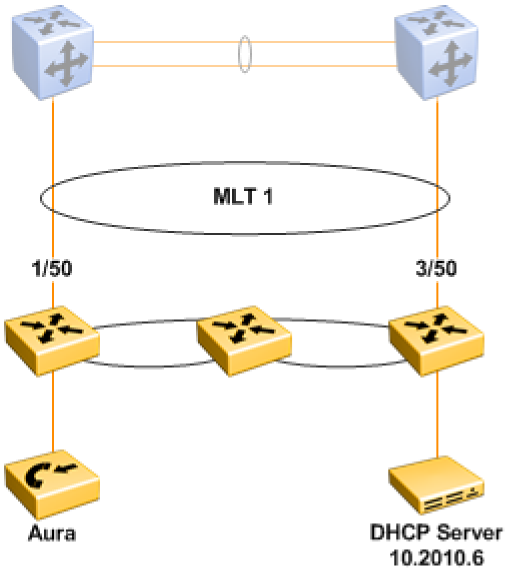

Figure 3.6 – ERS 5000 Data Center Switch Configuration

3.6.1 Virtual LANs

For this configuration step the following VLAN and IP parameters will be enabled:

ERS 5000	
VLAN Name	**VLAN ID**
Management	2000
DataCenter	2010

Table 3.6.1 – ERS 5000 VLAN Parameters

1 Create VLANs:

```
ERS5000-1(config)# vlan create 2000 name Management type port
ERS5000-1(config)# vlan create 2010 name DataCenter1 type port
```

2 Configure VLAN Membership and Port Tagging:

```
ERS2500-1(config)# vlan ports 1/50,3/50 tagging tagall filter-untagged-frame enable
ERS2500-1(config)# vlan members remove 1 all
```

```
ERS2500-1(config)# vlan members add 2000 1/50,3/50
ERS2500-1(config)# vlan members add 2010 1/1-3/50
```

3.6.2 Multilink Trunking (MLT)

For this configuration step the following MLT parameters will be enabled:

ERS 5000		
MLT ID	**Port(s)**	**Learning**
1	1/50,3/50	Disabled

Table 3.6.2 – ERS 5000 MLT Parameters

1 Create MLT:

```
ERS5000-1(config)# mlt 1 name Dist-CD member 1/50,3/50 learning disabled
ERS5000-1(config)# mlt 1 loadbalance advance
ERS5000-1(config)# mlt 1 enable
```

Tip – When configuring the MLT group, Spanning Tree must be disabled on the MLT. The *learning disable* on the MLT command disables Spanning Tree for the MLT.

Note – In an RSMLT square/full mesh core topology, Avaya recommends to set the MLT algorithm on Avaya edge stackable switches to advance for IP based traffic. Tests have shown that this will increase performance over the default MLT setting of basic.

3.6.3 Management

For this configuration step the following Management parameters will be enabled:

ERS 5000		
VLAN Name	**VLAN ID**	**IP Address**
Management	2000	10.20.0.17/24

Table 3.6.3 – ERS 5000 Management Parameters

1 Specify a Management VLAN:

```
ERS5000-1(config)# vlan mgmt 2000
```

2 Create a Management IP Address:

```
ERS5000-1(config)# interface vlan 2000
ERS5000-1(config-if)# ip address 10.20.0.17 255.255.255.0
```

```
ERS5000-1(config-if)# exit
```

3.6.4 VLACP

For this configuration step the following VLACP parameters will be enabled:

ERS 5000				
VLACP MAC Address	**Port(s)**	**Timeout**	**Fast Periodic Time**	**Timeout Scale**
01:80:c2:00:00:0f	1/50,3/50	Short	500	5

Table 3.6.4 – ERS 5000 VLACP Parameters

1 Globally enable VLACP:

```
ERS5000-1(config)# vlacp macaddress 01:80:c2:00:00:0f
ERS5000-1(config)# enable
```

2 Enable VLACP on the MLT 1 ports:

```
ERS5000-1(config)# interface fastEthernet all
ERS5000-1(config-if)# vlacp port 1/50,3/50 timeout short
ERS5000-1(config-if)# vlacp port 1/50,3/50 fast-periodic-time 500
ERS5000-1(config-if)# vlacp port 1/50,3/50 timeout-scale 5
ERS5000-1(config-if)# vlacp port 1/50,3/50 enable
ERS5000-1(config-if)# exit
```

3.6.5 Spanning Tree Protocol / BPDU Filtering

For this configuration step the following Spanning Tree Protocol parameters will be enabled:

ERS 5000			
Port(s)	**Learning Mode**	**BPDU Filtering Timeout**	**BPDU Filtering Enable**
1/1-48,2/1-48,3/1-48	Fast	0	Enable

Table 3.6.5 – ERS 5000 Spaning Tree Protocol Parameters

△ Caution – Spanning Tree Protocol and BPDU filtering should only be enabled on edge ports and not MLT ports.

1	Enable Fast Start Spanning Tree Protocol and BPDU Filtering on Edge Ports:

```
ERS5000-1(config)# interface fastEthernet 1/1-48,2/1-48,3/1-48
ERS5000-1(config-if)# spanning-tree learning fast
ERS5000-1(config-if)# spanning-tree bpdu-filtering timeout 0
ERS5000-1(config-if)# spanning-tree bpdu-filtering enable
ERS5000-1(config-if)# exit
```

3.6.6 Rate Limiting

For this configuration step the following Rale Limiting parameters will be enabled:

ERS 5000	
Port(s)	**Broadcast / Multicast Rate (Percent)**
1/1-48,2/1-48,3/1-48	10

Table 3.6.6 – ERS 5000 Spaning Tree Protocol Parameters

1	Enable Broadcast and Multicast Rate Limiting on Edge ports:

```
ERS5000-1(config)# interface fastEthernet all
ERS5000-1(config-if)# rate-limit port 1/1-48,2/1-48,3/1-48 both 10
ERS5000-1(config-if)# exit
```

3.6.7 IGMP Snooping and Proxy

For this configuration step the following IGMP parameters will be enabled:

ERS 2500		
VLAN ID	**IGMP Snooping**	**IGMP Proxy**
2010	Yes	Yes

Table 3.6.7 – ERS 5000 IGMP Parameters

1	Enable IGMP Snooping and Proxy on User VLANs:

```
ERS2500-1(config)# vlan igmp 2010 snooping enable
ERS2500-1(config)# vlan igmp 2010 proxy enable
```

3.7 Key Health Indicators Configuration

For VSP 9000, Key Health Indicators (KHI) is always on and currently only supports Forwarding and Performance statistics. To configure KHI for ERS 8800, you enable and disable the feature globally or only some of the KHI types. Controlling KHI types has a greater impact on loaded systems.

The main configuration actions for KHI are:

- Enabling or disabling KHI (at global or feature-level)
- Displaying statistics
- Clearing statistics/history to establish a new timeline

Note – EDM does not support KHI configuration.

3.7.1 Displaying Key Health Indicator Information

1 Display all Key Health Indicator information:

```
ERS8800:5# config sys set khi info
```

3.7.2 Globally Enabling Key Health Indicators

The ERS 8800 provides a global boot delay parameter for KHI. If the system begins collecting statistics at boot-up, the transitions that the system initially experiences do not provide an appropriate baseline of normal. To provide a valid baseline, you can configure the **boot-delay** parameter to specify how long the system takes to stabilize before KHI begins collecting statistics.

1 Globally enable Key Health Indicators:

```
ERS8800:5# config sys set khi khi-enable true
ERS8800:5# config sys set khi boot-delay 5
```

3.7.3 Enable Management Key Health Indicators

Management Key Health Indicators tracks TCP connections, CLI users, and KHI log status.

1 Enable Management Key Health Indicators:

```
ERS8800:5# config sys set khi mgmt-khi-enable true
```

3.7.4 Enable Chassis Key Health Indicators

Chassis Key Health Indicators displays the chassis key health indicators, such as temperature, fans, power supply, slots and CPU state.

1 Enable Chassis Key Health Indicators:

```
ERS8800:5# config sys set khi chassis-khi-enable true
```

3.7.5 Enable Performance Key Health Indicators

Performance Key Health Indicators displays the performance key health indicators, such as utilization status for CPU and switch fabric.

1 Enable Performance Key Health Indicators:

```
ERS8800:5# config sys set khi performance-khi-enable true
```

3.7.6 Protocol Key Health Indicators

Protocol Key Health Indicators are always on and tracks the health of the following protocols:

- OSPF
- BGP
- IST/SMLT
- PIM
- IGMP
- VLACP
- RTM and FDB table statistics

Protocol Key Health Indicators also provides statistics and historical data for protocol and neighbor state transitions. It also allows for the establishment of reference timestamps and reference data to track protocol health in the network. It supports VRFs.

Every protocol has a large number of parameters that can be tracked, but only the key parameters are tracked by the Key Health Indicators. Protocol information is collected and displayed on-demand, creating minimal overhead. The information is not stored in any separate database (except reference data), so that memory utilization is also minimal.

To ensure the validity of the Key Health Indicators information, ensure that it is in sync with the output from the protocol show commands, and verify that the timestamps are relevant.

3.7.7 Enabling Forwarding Key Health Indicators

Forwarding Key Health Indicators provides a history of the last 10 Forwarding KHI events and tracks the following on each chassis slot:

- Asic Resets
- RSP State Error Events
- RSP Stats Error Events
- F2X (F2I, F2E) Error Events

1 Enable Forwarding Key Health Indicators:

```
ERS8800:5# config sys set khi forwarding-khi-enable true
```

3.7.8 IP Interface Key Health Indicators

IP Interface Key Health Indicators are always on and provides the total configured and total operational IP interface count. It also provides a history of the last 10 IP Interface Up/Down events.

3.7.9 Configure Port Key Health Indicators

Port Key Health Indicators track the following information:

- Overall system statistics (unicast, multicast and broadcast Rx, Tx packets) for the preceding 2 minutes
- Port Up/Down Events
- SMLT Port Up/Down Events
- IST Port Up/Down Events
- Port Errors

1 Enable Port Key Health Indicators:

```
ERS8800:5# config sys set khi port-khi-enable true
```

4. Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

4.1 Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

4.2 Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

4.3 Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

4.4 Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.