DEPARTMENT OF VETERANS AFFAIRS

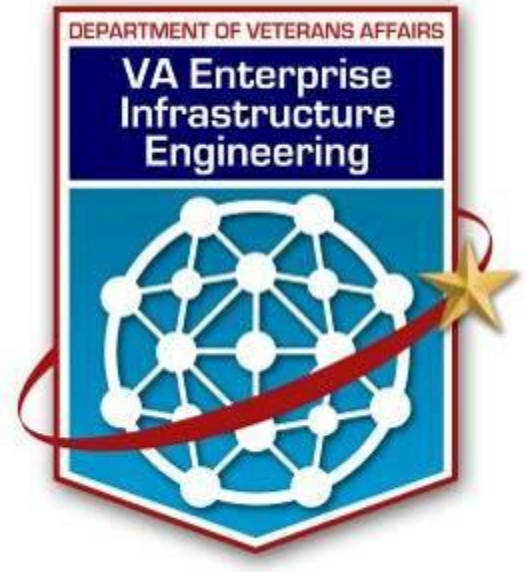

OFFICE OF INFORMATION AND TECHNOLOGY

ENTERPRISE INFRASTRUCTURE ENGINEERING

VA Enterprise IT Infrastructure Standard

Local Area Network

Production 1.0

November 5, 2009

TABLE OF CONTENTS

1 INTRODUCTION

1.1 PURPOSE

A standard is a set of rules or requirements that are determined by a consensus opinion of subject matter experts and prescribe criteria for a product, process, test or procedure. The general benefits of a standard are quality, interchangeability of parts or systems, and consistency. Information Technology (IT) standards are based on business needs provided through or supported by IT Services. IT Services are designed to support business processes and are constructed from software, hardware and infrastructure components. Establishing and enforcing standards for the selection and configuration of these supporting components improves the maintainability, reliability and availability of IT Services within projected economic constraints in alignment with business needs.

This standards document lists the acceptable and recommended specifications for Local Area Network. Sections include standard specifications for subject components, decisions supporting the standard specifications, guidelines or recommendations for implementing the standard specifications, and supplemental factors to consider when evaluating subject components. Other supplementary documents will provide guidance on procuring components that meet the standard specifications, guidance on integrating them with existing components, and explanation of how the subject components fit into the VA Architecture.

1.2 OBJECTIVES

This standard provides acceptable and recommended specifications to support:

- Solution Evaluation
- Requirement Evaluation
- Solution Design
- Solution Procurement and Bid Evaluation
- Evaluation of Architectural Specifications
- Click here to enter a list of objectives specific to this standard.

1.3 SCOPE

This standard applies to:

- This standard applies to the Region 2 and Region 3 Data Center Local Area Network

2 STANDARDS

2.1 CORE NETWORK SWITCH

This specification set identifies the requirements for a core layer switch.

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
1	Chassis	Type	Modular
		Rack Mounting	Required
		System Throughput	>= 1.44 Tbps
		Packet Processor Redundancy Support	Preferred, (See Implementation Guidance)
		Control Plane Redundancy Support	Required
		Front-Back Airflow Support	Required
		Cooling Redundancy Support	Required
		Power Supply Redundancy Support	Required
		Online Insertion/Removable Hot-Pluggable	Required
2	Packet Processor	Type	Modular/Upgradable
		Layer 3 Switching Throughput (pps)	>= 450 Mpps
		Layer 2 Switching Throughput (pps)	>= 450 Mpps
		Online Insertion/Removal Support	Required
		Memory	>= 1 GB
		L3 Protocols switched in hardware	Required: IPv4, IPv6, MPLS / Preferred: GRE, NAT

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
3	Linecards/Port Interfaces	Interface Speed	10 Gigabit Ethernet and 1 Gigabit Ethernet
		Media Flexibility Support	MMF >= 26m, SMF >= 10km, SMF >= 40km, MMF >= 300m, MMF(850nm)-10GBASE-S, SMF(1310nm)-10GBASE-L, SMF(1550nm)-10GBASE-E, MMF/SMF-10GBASE-LX4, 10GBase-T(Cat6a)
		Port per ASIC (Oversubscription)	1:1 Preferred
		Throughput to Backplane	>= 40 Gbps
		Queuing Properties	200MB/port w/ Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
		Online Insertion/Removable Support	Required
		On-Card distributed packet forwarding support	Preferred
4	Operating System	Type	Non-Monolithic (Modular) Preferred
		Features Supported	802.1p, 802.1q, Broadcast Storm Control
		Multi-chassis Link Aggregation Support	Required
		Multi-chassis Management Interface Aggregation	Preferred
		Protocols Supported	Required: Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP
		Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred
5	Support	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
		Parts Replacement	<= 24 hours, preferred <= 4 hours
6	High Availability	Operating System Maturity	Preferred released and QA Tested for >= 6 Months

ID	Primary Attribute	Secondary Attribute	Specification
		Hardware Maturity	Available for purchase for >= 3 Years
		Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a “real world” enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under “real world” traffic loads.

2.2 END-OF-ROW ACCESS LAYER SWITCH (10/100/1000 HOST CONNECTIVITY)

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type	Modular (Multi-Slot)
		Rack Mounting	Required
		System Throughput	>=720 Mbps
		Packet Processor Redundancy Support	Required
		Control Plane Redundancy Support	Required
		Front-Back Airflow Support	Required
		Cooling Redundancy Support	Required
		Power Supply Redundancy Support	Required
		Online Insertion/Removable Hot-Pluggable	Required
2	Packet Processor	Type	Modular/Upgradable
		Layer 3 Switching Throughput (pps)	>= 450 Mpps
		Layer 2 Switching Throughput (pps)	>= 450 Mpps

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Online Insertion/Removable Support	Required
		Memory	>= 1 GB
		Protocols switched in hardware	IPv4, IPv6, MPLS, GRE, NAT (hardware-assisted)
3	Linecards/Port Interfaces	Interface Speed	10/100 Megabit Ethernet , 1 Gigabit Ethernet & 10 Gigabit Ethernet
		Media Flexibility Support	MMF >= 220m, SMF >= 10km, UTP Cat-5 MMF(850nm)-10GBASE-S, SMF(1310nm)-10GBASE-L, SMF(1550nm)-10GBASE-E, MMF/SMF-10GBASE-LX4
		Port per ASIC (Oversubscription	1:1 Preferred
		Throughput to Backplane	>= 40 Gbps
		Queuing Properties	TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
		Online Insertion/Removable Support	Required
		On-Card distributed packet forwarding support	Preferred
4	Operating System	Type	Non-Monolithic (Modular) Preferred
		Features Supported	802.1p, 802.1q, Broadcast Storm Control
		Multi-chassis Link Aggregation Support	Required
		Multi-chassis Management Interface Aggregation Support	Preferred
		Protocols Supported	Required: Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP
		Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred

ID	Primary Attribute	Secondary Attribute	Specification
5	Support	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
		Parts Replacement	<= 24 hours, preferred <= 4 hours
6	High Availability	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
		Hardware Maturity	Available for purchase for >= 3 Years
		Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a “real world” enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under “real world” traffic loads.

2.3 HIGH-BANDWIDTH CONVERGED I/O ACCESS LAYER SWITCH (10GE HOST CONNECTIVITY)

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type	Modular
		Rack Mounting	Required
		System Throughput	>= 1 Tbps
		Packet Processor Redundancy Support	Preferred, (See Evaluation Criteria)
		Control Plane Redundancy Support	Preferred, (See Evaluation Criteria)
		Front-Back Airflow Support	Required
		Cooling Redundancy Support	Required
		Power Supply Redundancy Support	Required
		Online Insertion/Removable	Required

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Hot-Pluggable	
2	Packet Processor	Type	Integrated
		Layer 3 Switching Throughput (pps)	>= 700 Mpps
		Layer 2 Switching Throughput (pps)	>= 700 Mpps
		Online Insertion/Removal Support	Required when modular
		Memory	>= 1 GB
		L3 Protocols switched in hardware	Preferred: IPv4, IPv6, MPLS, GRE, NAT (See Evaluation Criteria)
3	Linecards/Port Interfaces	Interface Speed	10 Gigabit Ethernet and 1 Gigabit Ethernet
		Media Flexibility Support	Multi-mode fiber, Single-mode fiber, Copper(SFP+, Twin-Axial, UTP)
		Port per ASIC (Oversubscription)	1:1 Preferred
		Throughput to Backplane	n/a
		Queuing Properties	200MB/port w/ Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
		Online Insertion/Removable Support	Required
		On-Card distributed packet forwarding support	n/a
4	Operating System	Type	Non-Monolithic (Modular) Preferred
		Features Supported	802.1p, 802.1q, Broadcast Storm Control
		Multi-chassis Link Aggregation Support	Required
		Multi-chassis Management	Preferred, (See Implementation Guidance)

ID	Primary Attribute	Secondary Attribute	Specification
		Interface Aggregation	
		Protocols Supported	Required: Jumbo Frames, Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP
		Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred
5	Support	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a network down scenario.
		Parts Replacement	<= 24 hours, preferred <= 4 hours, (See Evaluation Criteria)
6	High Availability	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
		Hardware Maturity	Available for purchase for >= 3 Years
		Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a “real world” enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under “real world” traffic loads.

2.4 VISTA PRIVATE SCS DCI ROUTER/SWITCH - 10MBPS

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type	Modular
		Rack Mounting	Required
		System Throughput	>= 20,000pps (10Mbps)
		Packet Processor Redundancy Support	Not required
		Control Plane Redundancy Support	Not required
		Front-Back Airflow Support	Required

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Cooling Redundancy Support	Required
		Power Supply Redundancy Support	Required
		Online Insertion/Removable Hot-Pluggable	Required
2	Packet Processor	Type	Fixed Configuration
		Layer 3 Switching Throughput (pps)	>= 20,000pps (10Mbps)
		Layer 2 Switching Throughput (pps)	N/A
		Online Insertion/Removal Support	Not required
		Memory	>= 256 MB DRAM, >= 64 MB Flash
		L3 Protocols switched in hardware	Required: IPv4, IPv6, MPLS / Preferred: GRE, NAT
3	Linecards/Port Interfaces	Interface Speed	10/100Mbps Ethernet, DS3 Serial
		Media Flexibility Support	Not required
		Port per ASIC (Oversubscription)	<= 8:1 Preferred
		Throughput to Backplane	>= 1 Gbps
		Queuing Properties	Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
		Online Insertion/Removable Support	Required
		On-Card distributed packet forwarding support	N/A
4	Operating System	Type	Non-Monolithic (Modular) Preferred

ID	Primary Attribute	Secondary Attribute	Specification
		Features Supported	802.1p, 802.1q, Broadcast Storm Control
		Multi-chassis Link Aggregation Support	Not required
		Multi-chassis Management Interface Aggregation	Not required
		Protocols Supported	Required: Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP
		Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred
5	Support	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a "network down" scenario.
		Parts Replacement	<= 24 hours, preferred <= 4 hours
6	High Availability	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
		Hardware Maturity	Available for purchase for >= 3 Years
		Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a "real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

2.5 BLADE SEVER CHASSIS ACCESS LAYER SWITCH (10/100/1000 HOST CONNECTIVITY)

Assumptions: High Availability (Yes/No?) Redundant packet processing? Redundany power supplies?

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type	Fixed, 2 per Blade Server
		Rack Mounting	N/A
		System Throughput	Up to 128Gbps
		Packet Processor	Not required in single switch. Provided by redundant

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Redundancy Support	switches in same blade server
		Control Plane Redundancy Support	Not required
		Front-Back Airflow Support	Not required
		Cooling Redundancy Support	Not required
		Power Supply Redundancy Support	Not required
2	Packet Processor	Type	Modular or non-modular, usually a fixed configuration choice
		Layer 3 Switching Throughput (pps)	35.5Mpps
		Layer 2 Switching Throughput (pps)	35.5Mpps
		Memory	256MB or greater
		Protocols switched in hardware	IPv4, IPv6
3	Interfaces	Interface Speed	1Gb required, 10Gb optional
		ASIC per port	1:1 Preferred
		Throughput to Backplane	N/A
		Queuing Properties	Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
		Online Insertion/Removable Hot-Pluggable Support	N/A
		On-Card distributed packet forwarding support	N/A
4	Operating System	Type	Modular or Monolithic
		Features Supported	802.1p, 802.1q, Broadcast Storm Control

ID	Primary Attribute	Secondary Attribute	Specification
		Multi-chassis Link Aggregation Support	Required
		Multi-chassis Management Interface Aggregation	Preferred
		Protocols Supported	Required: Jumbo Frames, Link-Layer Discover Protocol, / Preferred: CDP, EIGRP
		Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred
5	Supportability	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
		Parts Replacement	<= 24 hours, preferred <= 4 hours
6	High Availability	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
		Hardware Maturity	Available for purchase for >= 3 Years
		Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a “real world” enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under “real world” traffic loads.

2.6 MANAGEMENT ACCESS LAYER SWITCH

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type	Fixed configuration, Stackable preferred
		Rack Mounting	Required
		System Throughput	24 to 32Gps
		Packet Processor	Not Required

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Redundancy Support	
		Control Plane Redundancy Support	Not Required
		Front-Back Airflow Support	Preferred
		Cooling Redundancy Support	Not Required
		Power Supply Redundancy Support	Preferred, may be supplied with external power supply.
		Online Insertion/Removable Hot-Pluggable	Not Required
2	Packet Processor	Type	Fixed Configuration
		Layer 3 Switching Throughput (pps)	None
		Layer 2 Switching Throughput (pps)	>=13 Mpps
		Online Insertion/Removable Support	Not Required
		Memory	>= 32MB
		Protocols switched in hardware	Layer 2 only, IPv4 and IPv6 packets
3	Linecards/Port Interfaces	Interface Speed	10/100 Megabit Ethernet, 1Gps for uplink interfaces
		Media Flexibility Support	Copper twisted pair with modular type uplink interfaces (copper or fiber). Other interface type may be required in an ad hoc basis.
		Port per ASIC (Oversubscription)	Permitted , <=6:1
		Throughput to Backplane	24 to 32 Gps
		Queuing Properties	TX: >= 4 RX: >= 4
		Online Insertion/Removable	Not Required

ID	Primary Attribute	Secondary Attribute	Specification
		Support	
		On-Card distributed packet forwarding support	Preferred
4	Operating System	Type	Monolithic or Modular
		Features Supported	802.1p, 802.1q, Broadcast Storm Control
		Multi-chassis Link Aggregation Support	Not Required
		Multi-chassis Management Interface Aggregation Support	Not Required , ability to Stack Preferred
		Protocols Supported	Required: LLDP, IPv4 and IPv6 packet support, Preferred: CDP, VTP
		Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred
5	Support	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
		Parts Replacement	<= 24 hours, preferred <= 4 hours
6	High Availability	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
		Hardware Maturity	Available for purchase for >= 3 Years
		Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a “real world” enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under “real world” traffic loads.

2.7 LOAD BALANCER

ID	Secondary Attribute	Specification

1	Rack Mountable	Required
2	Load Balancing Predictors	Least loaded, Least Bandwidth, Least Connections, Round-Robin, Hash Address
3	Server Health Monitoring	ICMP, TCP, ECHO(tcp&udp), HTTP, HTTPs, FTP, Telnet, DNS, RADIUS, SNMP, Scripted
4	Redundancy	Required
5	Role-Based Administration	Required
6	Throughput	>= 4 Gbps
7	Packets per Second	>= 6.5 million
8	NAT Entries	>= 1 million
9	Connections per Second	>= 325,000 sustained
10	Concurrent Connections	>= 4 million

3 SUPPORTING DETAILS FOR STANDARDS

3.1 CORE NETWORK SWITCH

It is recognized that the local area network industry is at a juncture in which two competing options could be explored in identifying an appropriate Core switch for the National Data Centers. The increasing popularity of server virtualization, blade server chassis, network attached storage (NAS), and the introduction of the convergence of storage traffic and traditional data traffic onto the LAN has increased the potential need to support higher density 10 gigabit Ethernet.

The current de-facto standard for the current VA Data Centers exploits the Catalyst V-E-Series 6500 as the Core or Aggregation switch. The Catalyst 6500 is a serial cross-bar switch with a total aggregate throughput of 1.44 Terabits per second when configured in a Virtual Switching System (VSS) Model. Cisco has announced plans for a next-generation supervisory module that will double the per-slot throughput to 80 Gigabit increasing the aggregate system throughput to 2.88 Terabits per second in a VSS configuration.

To satisfy high-density 10 gigabit Ethernet requirements at the aggregation and core layers that exceed the capabilities of the Catalyst platform Cisco has introduced the Nexus 7000 which introduces the capability of an aggregate throughput greater than 15 Terabits per second. The Nexus 7000 was announced and started shipping the second quarter of 2008 and is Cisco's next generation Data Center switch and is targeted to replace the Catalyst 6500 in the Data Center.

It has been communicated to the authors of this document that it is the position of the Executive CTO for VA Enterprise Infrastructure Engineering and the Senior Engineering Group that the proliferation of converged I/O and high-densities of 10Gigabit Ethernet in the next three to five years is not imminent. It is therefore the position of this document that the current de-facto standard solution will continue to meet the needs of the organization and therefore should continue to be procured and implemented. When application requirements change necessitating the implementation of a core switch with higher aggregate throughput capacities it will be possible to re-deploy the Catalyst from the core to an access layer.

3.1.1 CHASSIS

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Modular
	Rack Mounting	Required
	System Throughput	>= 1.44 Tbps
	Packet Processor Redundancy Support	Preferred, (See Implementation Guidance)
	Control Plane Redundancy Support	Required
	Front-Back Airflow Support	Required
	Cooling Redundancy Support	Required
	Power Supply Redundancy Support	Required
	Online Insertion/Removable Hot-Pluggable	Required

EXPLANATION OF STANDARD

When there are high density port count requirements that are supporting a mission critical application a modular highly redundant high availability switching architecture will be required. A modular chassis will allow for the upgrade of line-card packet processor modules and interface modules providing the flexibility to replace failed or upgrade obsolete components over time as new technology is introduced.

At the time of the writing of this document, specific data identifying the performance criteria for the applications supported by the Core switching platform was not available; therefore, the authors utilized switching performance data from existing architectures in place today. The system throughput of a switching system, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements of the VA as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second. There is no indication that even with a deployment of 10 Gigabit Ethernet attached virtual servers and storage systems that the VA data center requirements will begin to approach even 100 Gigabits per second of aggregate throughput.

A redundant packet processor is an essential component to a high-availability data network which provides non-stop forwarding and state-full failover during a negative event. Control Plane Redundancy provides the ability to recover from a negative event targeting the router/switch control plane. Additionally a redundant control plane provides the ability to perform in-service software upgrades, take advantage of nonstop forwarding and state-full switchover to a redundant processor.

During normal operation, electrical components of a switching system product heat which is necessary to be dissipated to maintain proper operation. Many implementations of switching systems have a system protection feature that will shut down a system if it overheats. To satisfy a high availability system it is critical to include N+1 redundancy for both cooling and system power.

Front-Back Airflow is a desirable feature to a core network switch that will reside in a data center design and will allow for hot-cold aisle deployment and contribute to more efficient cooling and increased energy efficiency.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of “hot-pluggable” components.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Ability to perform In Service Software Upgrades.
- Slot count of Chassis and Port count of Modules
- Packet Processor Redundancy
- Packet Switching performance
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

There will be two switches purchased and configured into a Virtual Switching System (VSS) model. If possible each of the switches will be placed in different locations in the data center suite in an effort to mitigate the affect of a negative event within a zone from creating a catastrophic network failure.

Packet processor redundancy is preferred and can be accommodated with a single packet processor in each chassis of a multi-chassis design where the control and data planes are merged between two chassis as is accomplished in the Virtual Switching System (VSS).

As a data center is implemented a concerted effort should be made to ensure that the redundant core switches have physical separation to reduce the impact of an event within a zone of the data center having a global impact on the system. As the data center expands every effort should be made to continue to maintain as much physical separation of the core switches as is reasonably possible.

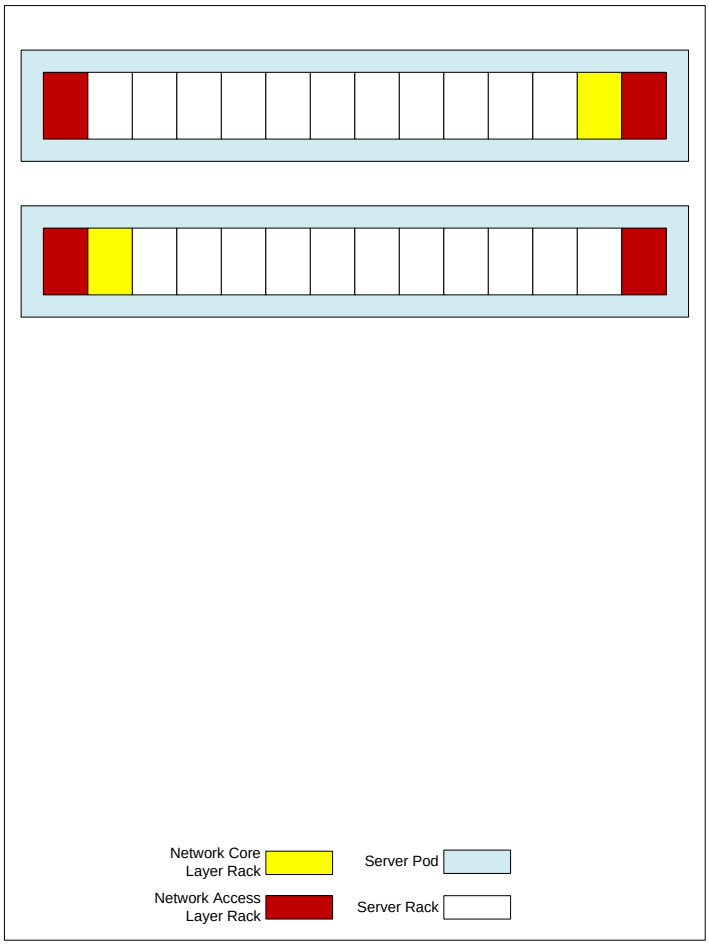

Figure 1 - Initial Implementation of a Data Center

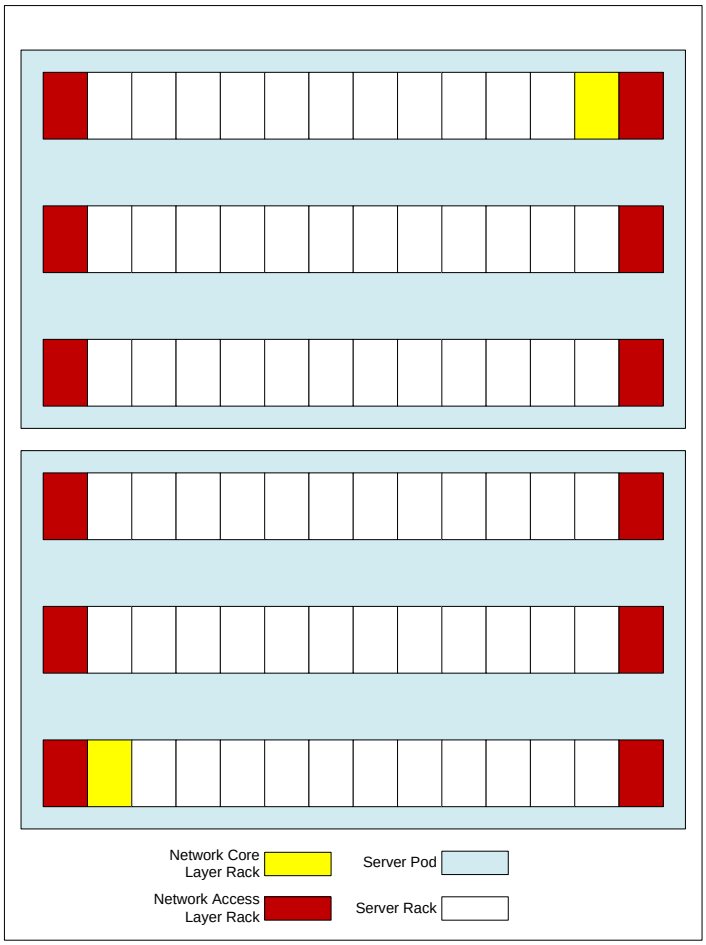

Figure 2 - Expansion of a Data Center

3.1.2 PACKET PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Type	Modular/Upgradable
	Layer 3 Switching Throughput (pps)	>= 450 Mpps
	Layer 2 Switching Throughput (pps)	>= 450 Mpps
	Online Insertion/Removal Support	Required
	Memory	>= 1 GB
	L3 Protocols switched in hardware	Required: IPv4, IPv6, MPLS / Preferred: GRE, NAT

EXPLANATION OF STANDARD

A modular and upgradable packet processor is desirable to allow for the leveraging of new technologies while protecting the investment in purchased interfaces and chassis.

At the time of the writing of this document, specific data identifying the performance criteria for the applications supported by the Core switching platform was not available; therefore, the authors utilized switching performance data from existing architectures in place today. The throughput of a packet processor, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second. There is no indication that even with a deployment of 10 Gigabit Ethernet attached virtual servers and storage systems that the VA data center requirements will begin to approach even 100 Gigabits per second of aggregate throughput.

The capability of a packet processor to allow of online insertion and removal is an essential component to a high-availability network design to allow for in-service technology upgrades and failed component replacement without interrupting critical traffic flows.

The memory of the packet processor translates to the capacity of layer 2 and layer 3 forwarding tables.

The Layer 3 protocols that can be switched in hardware will directly affect the performance of the entire system. A packet processor that can switch critical protocols such as IPv4, IPv6 and MPLS in hardware is a requirement.

EVALUATION FACTORS

- Required support for Multi-Chassis Switch Aggregation feature to support High Availability
- Ability to perform In Service Software Upgrades.
- Stability: Will be implemented in a Healthcare environment
- Packet Switching performance
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.1.3 LINECARDS/PORT INTERFACES

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Interface Speed	10 Gigabit Ethernet and 1 Gigabit Ethernet
	Media Flexibility Support	MMF >= 26m, SMF >= 10km, SMF >= 40km, MMF >= 300m, MMF(850nm)-10GBASE-S, SMF(1310nm)-10GBASE-L, SMF(1550nm)-10GBASE-E, MMF/SMF-10GBASE-LX4, 10GBase-T(Cat6a)
	Port per ASIC (Oversubscription)	1:1 is preferred, See Evaluation Criteria
	Throughput to Backplane	>= 40 Gbps
	Queuing Properties	200MB/port w/ Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
	Online Insertion/Removable Support	Required
	On-Card distributed	Preferred, See Evaluation Criteria

	packet forwarding support	

EXPLANATION OF STANDARD

Interface Speed- One Gigabit Ethernet Interface speeds will need to be met for Server connectivity. Ten Gigabit Ethernet speeds will need to be supported for Uplinks to Core

Linecards should be able to support a variety of physical media to allow for the flexibility of connecting different technologies over varying distances which can have different tolerances for that are dependent on the media used to transport the signals.

Ports per ASIC (oversubscription) – Depending on performance requirements oversubscription can be perfectly acceptable. As we move to supporting virtualization and I/O convergence resulting in integrating storage technologies onto the LAN the VA will be required to support very minimal oversubscription and therefore the standards will be adapted to leverage solutions and technologies that will support a non-blocking architecture.

The throughput capabilities of a linecard, while largely theoretical, does provide a loose scale by which one component can be measured against another. The standards outlined above demonstrate the current capabilities of deployed systems that are satisfying requirements of the VA now and for the foreseeable future.

Queuing - The queuing properties demonstrate the components ability to allow packets to be en-queued for transmission on a port at a rate greater than the physical medium can support. The ability for port interfaces to have compound buffers and queues translates directly to the performance on the overall system. With compound buffers and queues it becomes possible to classify different applications and assigning appropriate priorities to those traffic flows enabling the system to queue important and time sensitive traffic in front of less important and less time sensitive traffic.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of "hot-pluggable" components.

Distributed packet forwarding capabilities of linecards extend the intelligence of forwarding packets between ports on a card, or between similarly capable cards in a chassis directly without requiring the packet to be processed on the system packet processor. This eliminates a “bottle-neck” of traffic flows and allows for a significant performance increase to be realized.

EVALUATION FACTORS

- Ports per ASIC (oversubscription) – The fewer ports that are sharing an ASIC the better the performance. The ideal oversubscription ratio is 1:1.
- On-card distributed packet forwarding support improves performance and is preferred.

IMPLEMENTATION GUIDANCE

If addition hardware is required for full performance it should be required in the original proposal.

3.1.4 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	Non-Monolithic (Modular) Preferred
	Features Supported	802.1p, 802.1q, Broadcast Storm Control
	Multi-chassis Link Aggregation Support	Required
	Multi-chassis Management Interface Aggregation	Required
	Protocols Supported	Required: Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP (See Evaluation Criteria)
	Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred

EXPLANATION OF STANDARD

Type – A modular Operating System is preferred due to its ability to upgrade individual components of the OS. This reinforces the high availability design model.

Features Supported (layer 2) – Specification covers standard layer 2 features supporting high availability and Quality of Service.

Multi-chassis Link Aggregation – Specification is required to support a high availability architecture.

Multi-chassis Management Interface – Managing multiple chassis via a single unified management interface would minimize complexity and would be a preferred method to manage data center switches.

Protocols – Today we rely heavily on CDP to manage data center infrastructure which is Cisco proprietary. LLDP is the standards based Link-Layer Discover Protocol. As a tool to help continue managing the data center the specification requires some form of link layer discovery protocol.

Network Management support – In managing the LAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

- The CDP protocol is a useful and desirable tool when troubleshooting networks
- The VA uses EIGRP exclusively for VA Medical Centers and EIGRP is highly preferred to leverage existing expertise and maintain compatibility with existing deployments.

IMPLEMENTATION GUIDANCE

3.1.5 SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
	Parts Replacement	<= 24 hours, preferred <= 4 hours

EXPLANATION OF STANDARD

The Data Center hardware and software design needs to include the necessary vender support required to maintain the operational state of a Highly Available Data Center.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.1.6 HIGH AVAILABILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
	Hardware Maturity	Available for purchase for >= 3 Years
	Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a "real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

EXPLANATION OF STANDARD

The VA data network supports medical systems that are used in the day-to-day care of patients with applications including the hospital information system, ancillary clinical support systems such as pharmaceutical distribution and bio-diagnostic systems to name a few; all of which are used in the treatment of patients. The availability of the network can influence the information available to providers that are making life-or-death decisions and it therefore becomes critical that the network be available 24x7x365. To support the high availability environment it is essential that the maturity of both the hardware and the supporting software be subjected to a rigorous evaluation of a quality assurance program to mitigate the risk of un-scheduled system down time. The solution should be operating in a non-critical environment for a time sufficient enough to identify any flaws or vulnerabilities that may compromise the stability or security of the system. The more time the system has matured will be directly related to the probability that anomalies will be identified and eliminated and there is a corollary to reduced risk of an outage.

EVALUATION FACTORS

- A minimum of three years available for purchase is desired.

IMPLEMENTATION GUIDANCE

3.2 END-OF-ROW ACCESS LAYER SWITCH (10/100/1000 HOST CONNECTIVITY)

3.2.1 CHASSIS

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Modular (Multi-Slot)
	Rack Mounting	Required
	System Throughput	>=720 Mbps
	Packet Processor Redundancy Support	Required
	Control Plane Redundancy Support	Required
	Front-Back Airflow Support	Required
	Cooling Redundancy Support	Required
	Power Supply Redundancy Support	Required
	Online Insertion/Removable Hot-Pluggable	Required

EXPLANATION OF STANDARD

When there are high density port count requirements that are supporting a mission critical application a modular highly redundant high availability switching architecture will be required. A modular chassis will allow for the upgrade of line-card packet processor modules and interface modules providing the flexibility to replace failed or upgrade obsolete components over time as new technology is introduced.

At the time of the writing of this document, specific data identifying the performance criteria for the applications supported by the access switching platform was not available; therefore, the authors utilized switching performance data from existing architectures in place today. The system throughput of a switching system, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements of the VA as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second. There is no indication that even with a deployment of 10 Gigabit Ethernet

attached virtual servers and storage systems that the VA data center requirements will begin to approach even 100 Gigabits per second of aggregate throughput.

A redundant packet processor is an essential component to a high-availability data network which provides non-stop forwarding and state-full failover during a negative event. Control Plane Redundancy provides the ability to recover from a negative event targeting the router/switch control plane. Additionally a redundant control plane provides the ability to perform in-service software upgrades, take advantage of nonstop forwarding and state-full switchover to a redundant processor.

During normal operation, electrical components of a switching system product heat which is necessary to be dissipated to maintain proper operation. Many implementations of switching systems have a system protection feature that will shut down a system if it overheats. To satisfy a high availability system it is critical to include N+1 redundancy for both cooling and system power.

Front-Back Airflow is a desirable feature to a core network switch that will reside in a data center design and will allow for hot-cold aisle deployment and contribute to more efficient cooling and increased energy efficiency.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. To be able to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of “hot-pluggable” components.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Ability to perform In Service Software Upgrades.
- Slot count of Chassis and Port count of Modules
- Packet Processor Redundancy
- Packet Switching performance
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.2.2 PACKET PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Type	Modular/Upgradable
	Layer 3 Switching Throughput (pps)	>= 450 Mpps
	Layer 2 Switching Throughput (pps)	>= 450 Mpps
	Online Insertion/Removable Support	Required
	Memory	>= 1 GB
	Protocols switched in hardware	IPv4, IPv6, MPLS, GRE, NAT (hardware-assisted)

EXPLANATION OF STANDARD

A modular and upgradable packet processor is desirable to allow for the leveraging of new technologies while protecting the investment in purchased interfaces and chassis.

At the time of the writing of this document, specific data identifying the performance criteria for the applications supported by the Core switching platform was not available; therefore, the authors utilized switching performance data from existing architectures in place today. The throughput of a packet processor, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second. There is no indication that even with a deployment of 10 Gigabit Ethernet attached virtual servers and storage systems that the VA data center requirements will begin to approach even 100 Gigabits per second of aggregate throughput.

The capability of a packet processor to allow of online insertion and removal is an essential component to a high-availability network design to allow for in-service technology upgrades and failed component replacement without interrupting critical traffic flows.

The memory of the packet processor translates to the capacity of layer 2 and layer 3 forwarding tables.

The Layer 3 protocols that can be switched in hardware will directly affect the performance of the entire system. A packet processor that can switch critical protocols such as IPv4, IPv6 and MPLS in hardware is a requirement.

EVALUATION FACTORS

- Required support for Multi-Chassis Switch Aggregation feature to support High Availability
- Ability to perform In Service Software Upgrades.
- Stability: Will be implemented in a Healthcare environment
- Packet Switching performance
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.2.3 LINECARDS / PORT INTERFACES

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Interface Speed	10/100 Megabit Ethernet , 1 Gigabit Ethernet & 10 Gigabit Ethernet
	Media Flexibility Support	MMF >= 220m, SMF >= 10km, UTP Cat-5 MMF(850nm)-10GBASE-S, SMF(1310nm)-10GBASE-L, SMF(1550nm)-10GBASE-E, MMF/SMF-10GBASE-LX4
	Port per ASIC (Oversubscription	1:1 Preferred
	Throughput to Backplane	>= 40 Gbps
	Queuing Properties	TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
	Online Insertion/Removable Support	Required
	On-Card distributed packet forwarding support	Preferred

EXPLANATION OF STANDARD

Interface Speed- One Gigabit Ethernet Interface speeds will need to be met for Server connectivity. Ten Gigabit Ethernet speeds will need to be supported for uplinks to Core and to the counter-switch in a VSS configuration

Linecards should be able to support a variety of physical media to allow for the flexibility of connecting different technologies over varying distances which can have different tolerances for that are dependent on the media used to transport the signals.

Ports per ASIC (oversubscription) – Depending on performance requirements oversubscription can be perfectly acceptable. As we move to supporting virtualization and I/O convergence resulting in integrating storage technologies onto the LAN the VA will be required to support very minimal oversubscription and therefore the standards will be adapted to leverage solutions and technologies that will support a non-blocking architecture.

The throughput capabilities of a linecard, while largely theoretical, does provide a loose scale by which one component can be measured against another. The standards outlined above demonstrate the current capabilities of deployed systems that are satisfying requirements of the VA now and for the foreseeable future.

Queuing - The queuing properties demonstrate the components ability to allow packets to be en-queued for transmission on a port at a rate greater than the physical medium can support. The ability for port interfaces to have compound buffers and queues translates directly to the performance on the overall system. With compound buffers and queues it becomes possible to classify different applications and assigning appropriate priorities to those traffic flows enabling the system to queue important and time sensitive traffic in front of less important and less time sensitive traffic.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of "hot-pluggable" components.

Distributed packet forwarding capabilities of linecards extend the intelligence of forwarding packets between ports on a card, or between similarly capable cards in a chassis directly without requiring the packet to be processed on the system packet processor. This eliminates a “bottle-neck” of traffic flows and allows for a significant performance increase to be realized.

EVALUATION FACTORS

- Ports per ASIC (oversubscription) – The fewer ports that are sharing an ASIC the better the performance. The ideal oversubscription ratio is 1:1.
- On-card distributed packet forwarding support improves performance and is preferred.

IMPLEMENTATION GUIDANCE

3.2.4 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	Non-Monolithic (Modular) Preferred
	Features Supported	802.1p, 802.1q, Broadcast Storm Control
	Multi-chassis Link Aggregation Support	Required
	Multi-chassis Management Interface Aggregation Support	Preferred
	Protocols Supported	Required: Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP
	Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred, and must be ODBC/JDBC complaint for future integration to the Enterprise Management Framework (EMF).

EXPLANATION OF STANDARD

Type – A modular Operating System is preferred due to its ability to upgrade individual components of the OS. This reinforces the high availability design model

Features Supported (layer 2) – Specification covers standard layer 2 features supporting high availability and Quality of Service.

Multi-chassis Link Aggregation – Specification is required to support a high availability architecture.

Multi-chassis Management Interface – Managing multiple chassis via a single unified management interface would be a preferred method in managing data center switches with simplicity.

Protocols – Today we rely heavily on CDP to manage data center infrastructure which is Cisco proprietary. LLDP is the standards based Link-Layer Discover Protocol. As a tool to help continue managing the data center the specification requires some form of link layer discovery protocol.

Network Management support – In managing the LAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

- The CDP protocol is a useful and desirable tool when troubleshooting networks

- The VA uses EIGRP exclusively for VA Medical Centers and EIGRP is highly preferred to leverage existing expertise and maintain compatibility with existing deployments.
-

IMPLEMENTATION GUIDANCE

3.2.5 SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
	Parts Replacement	<= 24 hours, preferred <= 4 hours

EXPLANATION OF STANDARD

The Data Center hardware and software design needs to include the necessary vender support required to maintain the operational state of a Highly Available Data Center.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.2.6 HIGH AVAILABILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
	Hardware Maturity	Available for purchase for >= 3 Years
	Quality Assurance	Employs a rigorous testing program in a laboratory environment that simulates a

	Certification	"real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

EXPLANATION OF STANDARD

The VA data network supports medical systems that are used in the day-to-day care of patients with applications including the hospital information system, ancillary clinical support systems such as pharmaceutical distribution and bio-diagnostic systems to name a few; all of which are used in the treatment of patients. The availability of the network can influence the information available to providers that are making life-or-death decisions and it therefore becomes critical that the network be available 24x7x365. To support the high availability environment it is essential that the maturity of both the hardware and the supporting software be subjected to a rigorous evaluation of a quality assurance program to mitigate the risk of un-scheduled system down time. The solution should be operating in a non-critical environment for a time sufficient enough to identify any flaws or vulnerabilities that may compromise the stability or security of the system. The more time the system has matured will be directly related to the probability that anomalies will be identified and eliminated and there is a corollary to reduced risk of an outage.

EVALUATION FACTORS

- A minimum of three years available for purchase is desired.
-

IMPLEMENTATION GUIDANCE

3.3 HIGH-BANDWIDTH CONVERGED I/O ACCESS LAYER SWITCH (10GE HOST CONNECTIVITY)

The data center of today is at a juncture in which an evolution of network protocol and physical transport for storage traffic and traditional Ethernet for data are converged onto one media. The increasing popularity of server virtualization, blade server chassis, network attached storage (NAS), and the introduction of the convergence of storage traffic and traditional data traffic onto the LAN has increased the potential need to support higher density 10 gigabit Ethernet.

To satisfy high-density 10 gigabit Ethernet requirements and consolidated I/O such as Fiber Channel over Ethernet (FCoE) in supporting ancillary systems that can be considered "less-than mission critical" it will be advantageous to implement "cutting edge" technology.

3.3.1 CHASSIS

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Static/Modular
	System Throughput	>= 1 Tbps
	Packet Processor Redundancy Support	Preferred, (See Evaluation Criteria)
	Control Plane Redundancy Support	Preferred, (See Evaluation Criteria)
	Front-Back Airflow Support	Preferred, (See Evaluation Criteria)
	Cooling Redundancy Support	Required
	Power Supply Redundancy Support	Required
	Online Insertion/Removable Hot-Pluggable	Required

EXPLANATION OF STANDARD

When there is high-bandwidth or converged I/O (FCoE) requirements a “state-of-the-art” high-capacity switching architecture will be required. A chassis that is static with modular slot for interfaces will allow for the upgrade of line-card interface modules providing the flexibility to replace failed or introduce new components over time.

The system throughput of a switching system, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements of the VA as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second. There is no indication that even with a deployment of 10 Gigabit Ethernet attached virtual servers and storage systems that the VA data center requirements will begin to approach even 100 Gigabits per second of aggregate throughput.

A redundant packet processor is an essential component to a high-availability data network which provides non-stop forwarding and state-full failover during a negative event. Control Plane Redundancy provides the ability to recover from a negative event targeting the router/switch control plane. Additionally a redundant control plane provides the ability to perform in-service software upgrades, take advantage of nonstop forwarding and state-full switchover to a redundant processor.

During normal operation, electrical components of a switching system product heat which is necessary to be dissipated to maintain proper operation. Many implementations of switching systems have a system protection feature that will shut down a system if it overheats. To satisfy a high availability system it is critical to include N+1 redundancy for both cooling and system power.

Front-Back Airflow is a desirable feature to a core network switch that will reside in a data center design and will allow for hot-cold aisle deployment and contribute to more efficient cooling and increased energy efficiency.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of “hot-pluggable” components.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Slot count of Chassis and Port count of Modules
- Packet Switching performance
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.3.2 PACKET PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Type	Preferred Modular/Upgradable, (See Evaluation Criteria)
	Layer 3 Switching Throughput (pps)	>= 700 Mpps
	Layer 2 Switching Throughput (pps)	>= 700 Mpps
	Online Insertion/Removal Support	Required when modular
	Memory	>= 512 MB
	L3 Protocols switched in hardware	n/a

EXPLANATION OF STANDARD

A modular and upgradable packet processor is desirable to allow for the leveraging of new technologies while protecting the investment in purchased interfaces and chassis.

At the time of the writing of this document, specific data identifying the performance criteria for the applications supported by the Core switching platform was not available; therefore, the authors utilized switching performance data from existing architectures in place today. The throughput of a packet processor, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second.

The capability of a packet processor to allow of online insertion and removal is an essential component to a high-availability network design to allow for in-service technology upgrades and failed component replacement without interrupting critical traffic flows.

The memory of the packet processor translates to the capacity of layer 2 and layer 3 forwarding tables.

The Layer 3 protocols that can be switched in hardware will directly affect the performance of the entire system. A packet processor that can switch critical protocols such as IPv4 and IPv6 in hardware is a requirement.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Packet Switching performance
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.3.3 LINECARDS/PORT INTERFACES

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Interface Speed	10 Gigabit Ethernet and 1 Gigabit Ethernet
	Media Flexibility Support	MMF >= 26m, SMF >= 10km, SMF >= 40km, MMF >= 300m, MMF(850nm)-10GBASE-S, SMF(1310nm)-10GBASE-L, SMF(1550nm)-10GBASE-E, MMF/SMF-10GBASE-LX4, 10GBase-T(Cat6a), TwinAxial(SFP+), 8/4/1 Gbps FiberChannel
	Port per ASIC (Oversubscription)	1:1 Preferred
	Throughput to Backplane	>= 375 Gbps
	Queuing Properties	200MB/port w/ Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
	Online Insertion/Removable Support	Required
	On-Card distributed packet forwarding support	Preferred

EXPLANATION OF STANDARD

Interface Speed- Ten Gigabit Ethernet Interface speeds with Fiber-Channel over Ethernet (FCoE) capabilities will need to be met for Server connectivity. Ten Gigabit Ethernet speeds will need to be supported for Uplinks to Core

Linecards should be able to support a variety of physical media and both the Ethernet and Fiber-Channel protocols to allow for the flexibility of connecting different technologies over varying distances which can have different tolerances for that are dependent on the media used to transport the signals.

Ports per ASIC (oversubscription) – Depending on performance requirements oversubscription can be perfectly acceptable. As we move to supporting virtualization and I/O convergence resulting in integrating storage technologies onto the LAN the VA will be required to support very minimal oversubscription and therefore the standards will be adapted to leverage solutions and technologies that will support a non-blocking architecture.

Queuing - The queuing properties demonstrate the components ability to allow packets to be en-queued for transmission on a port at a rate greater than the physical medium can support. The ability for port interfaces to

have compound buffers and queues translates directly to the performance on the overall system. With compound buffers and queues it becomes possible to classify different applications and assigning appropriate priorities to those traffic flows enabling the system to queue important and time sensitive traffic in front of less important and less time sensitive traffic.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of "hot-pluggable" components.

Distributed packet forwarding capabilities of linecards extend the intelligence of forwarding packets between ports on a card, or between similarly capable cards in a chassis directly without requiring the packet to be processed on the system packet processor. This eliminates a “bottle-neck” of traffic flows and allows for a significant performance increase to be realized.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Packet Switching performance
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.3.4 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	Non-Monolithic (Modular) Preferred
	Features Supported	802.1p, 802.1q, Broadcast Storm Control
	Multi-chassis Link Aggregation Support	Required
	Multi-chassis Management Interface Aggregation	Preferred
	Protocols Supported	Required: Jumbo Frames, Link-Layer Discover Protocol / Preferred: CDP, EIGRP
	Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred, and must be ODBC/JDBC complaint for future integration to the Enterprise Management Framework (EMF).

EXPLANATION OF STANDARD

Type – A modular Operating System is preferred due to its ability to upgrade individual components of the OS. This reinforces the high availability design model.

Features Supported (layer 2) – Specification covers standard layer 2 features supporting high availability and Quality of Service.

Multi-chassis Link Aggregation – Specification is required to support a high availability architecture.

Multi-chassis Management Interface – Managing multiple chassis via a single unified management interface would minimize complexity and would be a preferred method to manage data center switches.

Protocols – Today we rely heavily on CDP to manage data center infrastructure which is Cisco proprietary. LLDP is the standards based Link-Layer Discover Protocol. As a tool to help continue managing the data center the specification requires some form of link layer discovery protocol.

Network Management support – In managing the LAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Packet Switching performance
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.3.5 SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
	Parts Replacement	<= 24 hours, preferred <= 4 hours

EXPLANATION OF STANDARD

The Data Center hardware and software design needs to include the necessary vender support required to maintain the operational state of a Highly Available Data Center.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.3.6 HIGH AVAILABILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
	Hardware Maturity	Available for purchase for >= 1 Years
	Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a "real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

EXPLANATION OF STANDARD

In certain circumstances where there are very high-bandwidth requirements it may be desirable to implement "cutting-edge" technology with the potential for increased performance. Depending on the scenario it may be acceptable to incur a modest amount of risk by implementing a less mature solution as a trade off to a more "state-of-the-art" solution.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.4 VISTA PRIVATE SCS DCI ROUTER/SWITCH - 10MBPS

The private network for VistA supports VMS System Communications Services (SCS). The implementation guidance from HP related to SCS is that there must be 10Mbps of bandwidth available with latency characteristics similar to that of Ethernet. It is imperative that the cluster communications via SCS are maintained to sustain the membership of each node in the cluster. If cluster communication to a node fails then the operation of the entire cluster and processing of the VistA database can be interrupted. Due to the critical nature of this environment it is desirable to have two completely separate and redundant "air-gapped" networks to support the VistA private communications with each database node having a connection to each of the two networks. Since the SCS protocol is extremely resilient and the nodes of the cluster are not prone to adverse effects as a result of an interruption to one of

the private networks as long as the second network is operational, it is acceptable to reduce the design criteria of the network hardware components since high-availability is being achieved in redundancy of the private networks.

3.4.1 CHASSIS

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Modular
	Rack Mounting	Required
	System Throughput	>= 20,000pps (10Mbps)
	Packet Processor Redundancy Support	Not required
	Control Plane Redundancy Support	Not required
	Front-Back Airflow Support	Required
	Cooling Redundancy Support	Required
	Power Supply Redundancy Support	Required
	Online Insertion/Removable Hot-Pluggable	Required

EXPLANATION OF STANDARD

A modular chassis will allow for the upgrade of line-card network interface modules providing the flexibility to replace failed or upgrade obsolete components over time as new technology is introduced.

The requirements for the VistA SCS private network are 10Mbps; therefore, the system throughput should be able to sustain 10Mbps at a minimum.

The high-availability for the VistA SCS private network is achieved in the redundant physical networks; packet processor and control plane redundancy are not required.

During normal operation, electrical components of a switching system produce heat which is necessary to be dissipated to maintain proper operation. Many implementations of switching systems have a system protection

feature that will shut down a system if it overheats. To satisfy a high availability system it is critical to include N+1 redundancy for both cooling and system power.

Front-Back Airflow is a desirable feature to a core network switch that will reside in a data center design and will allow for hot-cold aisle deployment and contribute to more efficient cooling and increased energy efficiency.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade to a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of "hot-pluggable" components.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

Two independent solutions are required; each vista database node has two private NICs, private NIC #1 and NIC #2 connect to private router/switch #1 and router/switch #2. Two dedicated communication circuits each >= 10Mbps connects Metro data center A and Metro data center B for each SCS private network. A completely separate Spanning-Tree domain is maintained for both Private network SCS #1 and Private network SCS #2.

Each database node will have two private NICs for SCS communications. SCS NIC #1 from each database node will connect to SCS router/switch #1 and SCS NIC #2 from each database node will connect to SCS router/switch #2. The two router/switches SCS #1 and SCS#2 should have sufficient physical separation so as to protect from an event within a zone impacting both switches. To accommodate SCS communication to a quorum site; the SCS router/switch #1 in Metro data center A will be connected to a port that is configured as a L2TPv3 pseudo wire passing the SCS packets through a tunnel routed on the layer 3 network over the regional and/or VA backbone WAN to a port at the quorum site where the SCS NIC #1 from the quorum node is attached. A similar configuration from the SCS switch #2 in Metro data center B, will be connected to a port that is configured as a L2TPv3 pseudo wire passing the SCS packets through a tunnel routed on the layer 3 network to a port at the quorum site where the SCS NIC #2 from the quorum node is attached.

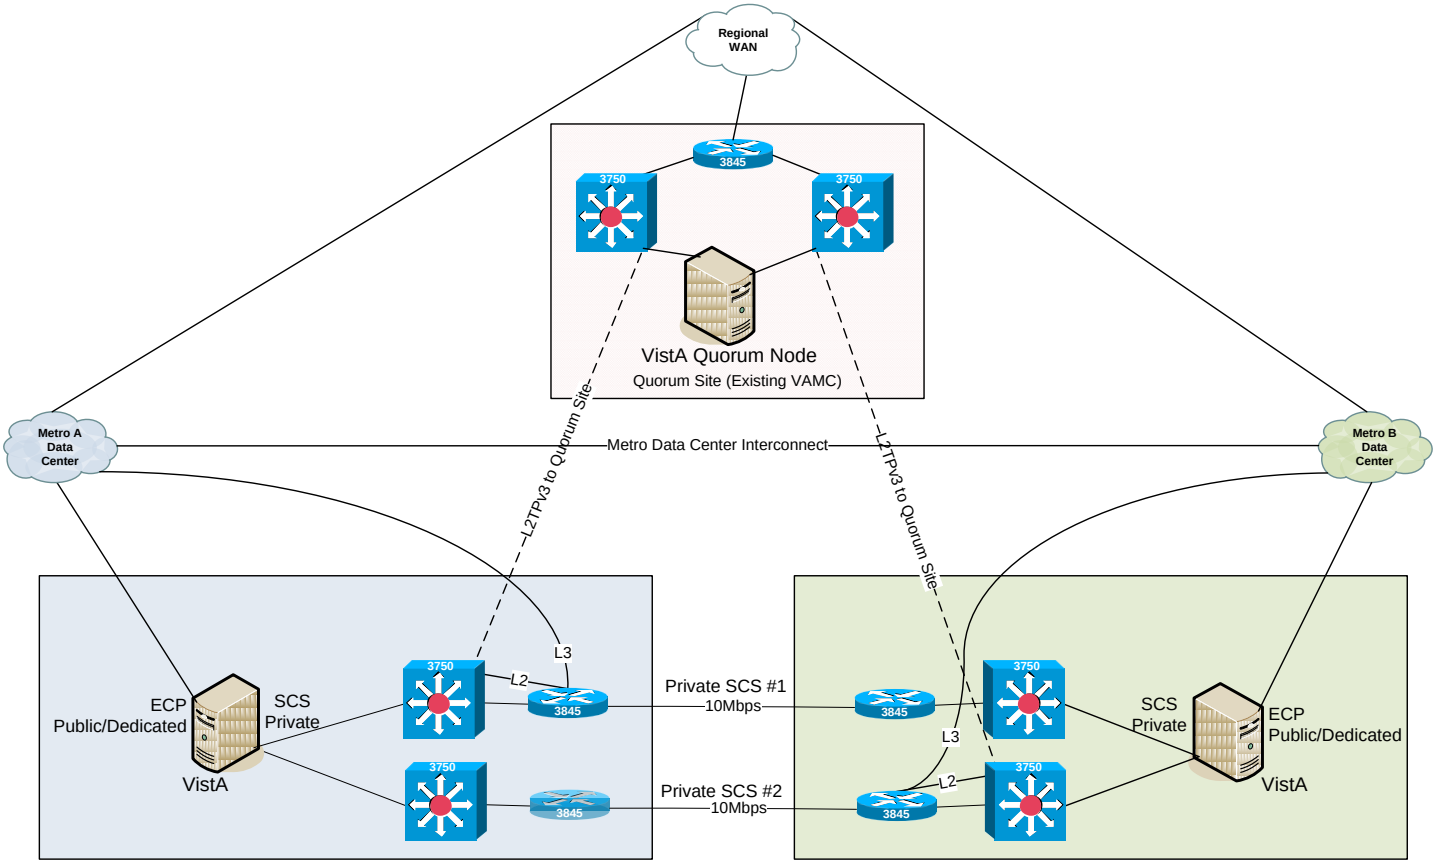

3.4.2 PACKET PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Type	Fixed Configuration
	Layer 3 Switching Throughput (pps)	>= 20,000pps (10Mbps)
	Layer 2 Switching Throughput (pps)	N/A

	Online Insertion/Removal Support	Not required
	Memory	>= 256 MB DRAM, >= 64 MB Flash
	L3 Protocols switched in hardware	Required: IPv4, IPv6, MPLS / Preferred: GRE, NAT

EXPLANATION OF STANDARD

The high-availability requirement for the VistA cluster communications is satisfied in the fully redundant physical network; therefore, a lower class router can be employed and a modular and upgradable packet processor is not required in this configuration.

The throughput of a packet processor, while largely theoretical, does provide a loose scale by which one system can be measured against another. The throughput requirements as identified by HP are 10Mbps.

With a static packet processor the capability of a packet processor to allow of online insertion and removal is not required.

The Layer 3 protocols that can be switched in hardware will directly affect the performance of the entire system. A packet processor that can switch critical protocols such as IPv4, IPv6 and MPLS in hardware is a requirement.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.4.3 LINECARDS/PORT INTERFACES

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Interface Speed	10/100Mbps Ethernet, DS3 Serial
	Media Flexibility Support	Not required
	Port per ASIC (Oversubscription)	<= 8:1 Preferred
	Throughput to Backplane	>= 1 Gbps
	Queuing Properties	Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4
	Online Insertion/Removable Support	Required
	On-Card distributed packet forwarding support	N/A

EXPLANATION OF STANDARD

Interface Speed- Per HP, 10Mbps Ethernet Interface speeds will be sufficient to meet VMS SCS connectivity requirements.

Linecards should be able to support a variety of physical media to allow for the flexibility of connecting different technologies over varying distances which can have different tolerances for that are dependent on the media used to transport the signals.

Ports per ASIC (oversubscription) – The performance requirements of the VMS SCS protocol are very modest and subsequently oversubscription is perfectly acceptable.

Queuing - The queuing properties demonstrate the components ability to allow packets to be en-queued for transmission on a port at a rate greater than the physical medium can support. The ability for port interfaces to have compound buffers and queues translates directly to the performance on the overall system. With compound buffers and queues it becomes possible to classify different applications and assigning appropriate priorities to

those traffic flows enabling the system to queue important and time sensitive traffic in front of less important and less time sensitive traffic.

Through the lifetime of a system it is occasionally necessary to replace failed components or to upgrade others with more memory or with a next generation component. The ability to replace such component while the system is fully operation with near zero interruption to the normal operation of the system is required in a high availability environment. To facilitate this requirement the chassis must have the capability to support online insertion and removal of "hot-pluggable" components.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.4.4 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	Non-Monolithic (Modular) Preferred
	Features Supported	802.1p, 802.1q, Broadcast Storm Control
	Multi-chassis Link Aggregation Support	Not required
	Multi-chassis Management Interface Aggregation	Not required
	Protocols Supported	Required: Link-Layer Discover Protocol, OSPF, BFD / Preferred: CDP, EIGRP
	Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred, and must be ODBC/JDBC complaint for future integration to the Enterprise Management Framework (EMF).

EXPLANATION OF STANDARD

Type – A modular Operating System is preferred due to its ability to upgrade individual components of the OS. This reinforces the high availability design model.

Features Supported (layer 2) – Specification covers standard layer 2 features supporting high availability and Quality of Service.

Protocols – Today we rely heavily on CDP to manage data center infrastructure which is Cisco proprietary. LLDP is the standards based Link-Layer Discover Protocol. As a tool to help continue managing the data center the specification requires some form of link layer discovery protocol.

Network Management support – In managing the LAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.4.5 SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
	Parts Replacement	<= 24 hours, preferred <= 4 hours

EXPLANATION OF STANDARD

The Data Center hardware and software design needs to include the necessary vender support required to maintain the operational state of a Highly Available Data Center.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.4.6 HIGH AVAILABILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
	Hardware Maturity	Available for purchase for >= 3 Years
	Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a "real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

EXPLANATION OF STANDARD

The VA data network supports medical systems that are used in the day-to-day care of patients with applications including the hospital information system, ancillary clinical support systems such as pharmaceutical distribution and bio-diagnostic systems to name a few; all of which are used in the treatment of patients. The availability of the network can influence the information available to providers that are making life-or-death decisions and it therefore becomes critical that the network be available 24x7x365. To support the high availability environment it is essential that the maturity of both the hardware and the supporting software be subjected to a rigorous evaluation of a quality assurance program to mitigate the risk of un-scheduled system down time. The solution should be operating in a non-critical environment for a time sufficient enough to identify any flaws or vulnerabilities that may compromise the stability or security of the system. The more time the system has matured will be directly related to the probability that anomalies will be identified and eliminated and there is a corollary to reduced risk of an outage.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.5 BLADE SEVER CHASSIS ACCESS LAYER SWITCH

3.5.1 CHASSIS

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Fixed, 2 per Blade Server
	Slot Count	N/A, fixed configuration (usually)
	Maximum Theoretical System Throughput	Up to 128Gbps
	Packet Processor Redundancy Support	Not required in single switch. Provided by redundant switches in same blade server
	Control Plane Redundancy Support	Not required
	Front-Back Airflow Support	Not required
	Cooling Redundancy Support	Not required
	Power Supply Redundancy Support	Not required

EXPLANATION OF STANDARD

Certain server platforms deployed by the VA will consist of Blade Server Chassis'. All Blade Servers typically contain two internal network switches to provide connectivity for the internal server blades. This standard defines the properties required for these switches. Most Blade Server manufacturers will provide a limited choice in switch type included with each server model. These choices are usually fixed configurations but may vary between server manufacturer.

Since the Blade Server switch implementations are usually offered as fixed configurations with very small form factors, many of the standard chassis requirements are either not applicable or can't be met by any manufacturer. Air flow and cooling are typically provided by the Blade Server fans themselves and not in the switches.

All Blade Server manufacturers should offer redundant internal switches as a standard option. This is a requirement. Typical Blade Server implementations utilize direct switch backplane connections to the individual blade servers and then multiple external 1Gig or 10Gig interfaces for uplinks to external switch environments.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.5.2 PACKET PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Modular or non-modular, usually a fixed configuration choice
	Layer 3 Switching Throughput (pps)	35.5Mpps
	Layer 2 Switching Throughput (pps)	35.5Mpps
	Multi-chassis Link Aggregation Support	Required
	Memory	256MB or greater
	Protocols switched in hardware	IPv4, IPv6 preferred

EXPLANATION OF STANDARD

Multi-chassis Link Aggregation support is the most important feature in this section. The VA has set a requirement that Multi-chassis Link Aggregation will be used to connect all non-Vista servers within the datacenters. What this means for the Blade Server switches is that both internal server switches must be able to bond all uplink ports to a single aggregated port channel and then bond all backplane connections to a one server to a single aggregated port channel. Each server will need it's own aggregated port channel.

At the time of the writing of this document, specific data identifying the performance criteria for the applications supported by the Core switching platform was not available; therefore, the authors utilized switching performance data from existing architectures in place today. The throughput of a packet processor, while largely theoretical, does provide a loose scale by which one system can be measured against another. The current throughput requirements as demonstrated in the currently deployed regional data centers are between one and two Gigabits per second. There is no indication that even with a deployment of 10 Gigabit Ethernet attached virtual servers and storage systems that the VA data center requirements will begin to approach even 100 Gigabits per second of aggregate throughput.

The Layer 3 protocols that can be switched in hardware will directly affect the performance of the entire system. A packet processor that can switch critical protocols such as IPv4 and IPv6 in hardware is a preference.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment

- Packet Switching performance
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.5.3 INTERFACES

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Interface Speed	1Gb required, 10Gb optional
	Interface Count	4-8+ 1Gbps or 2+ 10Gbps
	ASIC per port	1:1 Preferred
	Throughput to Backplane	N/A
	Queuing Properties	Shaped Round Robin + DSCP Support TX: Priority: Preferred, Standard: >= 4 RX: Standard: >= 4

EXPLANATION OF STANDARD

Typical Blade Server switches utilize backplane connections for the direct internal blade server uplink ports. External ports are provided to uplink the Blade Server switches to existing external switch fabrics.

Ports per ASIC (oversubscription) – Depending on performance requirements oversubscription can be perfectly acceptable. As we move to supporting virtualization and I/O convergence resulting in integrating storage technologies onto the LAN the VA will be required to support very minimal oversubscription and therefore the standards will be adapted to leverage solutions and technologies that will support a non-blocking architecture.

Queuing - The queuing properties demonstrate the components ability to allow packets to be en-queued for transmission on a port at a rate greater than the physical medium can support. The ability for port interfaces to have compound buffers and queues translates directly to the performance on the overall system. With compound buffers and queues it becomes possible to classify different applications and assigning appropriate priorities to those traffic flows enabling the system to queue important and time sensitive traffic in front of less important and less time sensitive traffic.

Certain specific switch port protocol features are required in order provide continuity with existing configurations. The above list is based on current VA implementations and existing server software capabilityes (Virtual Environments).

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment

- Packet Switching performance
- Experience VA currently has with solution
- Interface counts: 4-8+ 1Gbps or 2+ 10Gbps
- Desirable Protocols include LACP and IGMP snooping

IMPLEMENTATION GUIDANCE

3.5.4 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Monolithic or Modular
	Features Supported	802.1p, 802.1q, Broadcast Storm Control
	Multi-chassis Link Aggregation	Required
	Multi-chassis Management Interface Aggregation	Preferred
	Protocols Supported	Required: Jumbo Frames, Link-Layer Discover Protocol, / Preferred: CDP, EIGRP
	Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred, and must be ODBC/JDBC complaint for future integration to the Enterprise Management Framework (EMF).

EXPLANATION OF STANDARD

Since typical Blade Server switch implementations are fixed configurations, the operating system type (modular vs. non-modular) isn't as critical as with other environments.

Features Supported (layer 2) – Specification covers standard layer 2 features supporting high availability and Quality of Service.

Protocols – Today we rely heavily on CDP to manage data center infrastructure which is Cisco proprietary. CDP is also provided as an option with VMWare operating systems and becomes very useful when interacting with ESX server environments. LLDP is the standards based Link-Layer Discover Protocol. As a tool to help continue managing the data center the specification requires some form of link layer discovery protocol.

EVALUATION FACTORS

- Stability: Will be implemented in a Healthcare environment
- LAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.5.5 SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
	Parts Replacement	<= 24 hours, preferred <= 4 hours

EXPLANATION OF STANDARD

The Data Center hardware and software design needs to include the necessary vender support required to maintain the operational state of a Highly Available Data Center.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.5.6 HIGH AVAILABILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
	Hardware Maturity	Available for purchase for >= 3 Years
	Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a "real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

EXPLANATION OF STANDARD

The VA data network supports medical systems that are used in the day-to-day care of patients with applications including the hospital information system, ancillary clinical support systems such as pharmaceutical distribution and bio-diagnostic systems to name a few; all of which are used in the treatment of patients. The availability of the network can influence the information available to providers that are making life-or-death decisions and it therefore becomes critical that the network be available 24x7x365. To support the high availability environment it is essential that the maturity of both the hardware and the supporting software be subjected to a rigorous evaluation of a quality assurance program to mitigate the risk of un-scheduled system down time. The solution should be operating in a non-critical environment for a time sufficient enough to identify any flaws or vulnerabilities that may compromise the stability or security of the system. The more time the system has matured will be directly related to the probability that anomalies will be identified and eliminated and there is a corollary to reduced risk of an outage.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.6 MANAGEMENT ACCESS LAYER SWITCH

Most servers come with a remote management interfaces (RMIs). These management interfaces vary in designation by manufacturer IGO, RILO, and RMOC are examples). They permit remote access to perform basic management functions and allow connectivity when the server is in an impaired state. As a normal business process managers access the servers via RDP. RMIs are therefore a redundant or alternate connection path.

RMIs have low to medium density port count and limited traffic requirements. Due to the limited density and traffic requirements and alternate access methods a fixed configuration, non-redundant switching architecture will be required. Stackable switches with redundant power supply options are preferred.

3.6.1 CHASSIS

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	Fixed configuration, Stackable preferred
	Rack Mounting	Required
	System Throughput	24 to 32Gps
	Packet Processor Redundancy Support	Not Required
	Control Plane Redundancy Support	Not Required
	Front-Back Airflow Support	Preferred
	Cooling Redundancy Support	Not Required
	Power Supply Redundancy Support	Preferred, may be supplied with external power supply.
	Online Insertion/Removable Hot-Pluggable	Not Required

EXPLANATION OF STANDARD

Front-Back Airflow is a desirable feature due to the hot-cold aisle VA data center design standard.

Due to RMI access being just one of many ways to access servers a minimal amount of RMI access downtime is acceptable.

EVALUATION FACTORS

- None

IMPLEMENTATION GUIDANCE

- None

3.6.2 PACKET PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Type	Fixed Configuration
	Layer 3 Switching Throughput (pps)	None
	Layer 2 Switching Throughput (pps)	>=13 Mpps
	Online Insertion/Removable Support	Not Required
	Memory	>= 32MB
	Protocols switched in hardware	Layer 2 only, IPv4 and IPv6 packets

EXPLANATION OF STANDARD

A fixed configuration packet processor is desirable due to the low bandwidth requirements and the availability of alternate connectivity.

All layer three (routing) functions will reside with higher level devices.

Online insertion and removal is not required as this is an alternate method of connectivity and short periods of downtime for this device is permitted.

EVALUATION FACTORS

- None

IMPLEMENTATION GUIDANCE

- None

3.6.3 PORT INTERFACES

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Interface Speed	10/100 Megabit Ethernet, 1Gps for uplink interfaces
	Media Flexibility Support	Copper twisted pair with modular type uplink interfaces (copper or fiber). Other interface type may be required in an ad hoc basis.
	Port per ASIC (Oversubscription)	Permitted , <=6:1
	Throughput to Backplane	24 to 32 Gps
	Queuing Properties	TX: >= 4 RX: >= 4
	Online Insertion/Removable Support	Not Required

EXPLANATION OF STANDARD

Interface Speed - Management access layer switches and RMIs have limited traffic requirements therefore only 10/100Mbs is required.

Medical Flexibility Support - Remote Management Interfaces are usually of the UTP type.

Ports per ASIC (oversubscription) - Due to low performance requirements oversubscription can is acceptable.

Throughput to Backplane - The throughput capabilities of a backplane, while largely theoretical, does provide a loose scale by which one component can be measured against another. The standards outlined above demonstrate the current capabilities of deployed systems that are satisfying requirements of the VA now and for the foreseeable future.

Queuing - The queuing properties demonstrate the components ability to allow packets to be en-queued for transmission on a port at a rate greater than the physical medium can support. The ability for port interfaces to have compound buffers and queues translates directly to the performance on the overall system. With compound buffers and queues it becomes possible to classify different applications and assigning appropriate priorities to those traffic flows enabling the system to queue important and time sensitive traffic in front of less important and less time sensitive traffic.

Online Insertion/Removable Support - Online insertion and removal is not required as this is an alternate method of connectivity and short periods of downtime for this device is permitted.

EVALUATION FACTORS

- None

IMPLEMENTATION GUIDANCE

- None

3.6.4 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	Monolithic or Modular
	Features Supported	802.1p, 802.1q, Broadcast Storm Control
	Multi-chassis Link Aggregation Support	Not Required
	Multi-chassis Management Interface Aggregation Support	Not Required , ability to Stack Preferred

	Protocols Supported	Required: LLDP, IPv4 and IPv6 packet support, Preferred: CDP, VTP
	Network Management	Centralized AAA w/ Role Based Authorization, Syslog SNMP v2c, v3 preferred, and must be ODBC/JDBC complaint for future integration to the Enterprise Management Framework (EMF).

EXPLANATION OF STANDARD

Type – A modular Operating System is usually not associated with the type of switch but is acceptable.

Features Supported (layer 2) – Specification covers standard layer 2 features supporting Quality of Service.

Multi-chassis Link Aggregation – Not required

Multi-chassis Management Interface – Managing multiple chassis via a single unified management interface would be a preferred method in managing data center switches with simplicity.

Protocols – Today we rely heavily on CDP to manage data center infrastructure which is Cisco proprietary. LLDP is the standards based Link-Layer Discover Protocol. As a tool to help continue managing the data center the specification requires some form of link layer discovery protocol.

Network Management support – In managing the LAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

- None

IMPLEMENTATION GUIDANCE

- None

3.6.5 SUPPORT

STANDARD

ID	***Secondary Attribute***	***Specification***
5	Technical Support	24x7 Telephone/Web with customer option to immediately escalate to a senior engineer in a “network down” scenario.
	Parts Replacement	<= 24 hours, preferred <= 4 hours

EXPLANATION OF STANDARD

The Data Center hardware and software design needs to include the necessary vender support required to maintain the operational state of a Highly Available Data Center.

EVALUATION FACTORS

- None

IMPLEMENTATION GUIDANCE

- None

3.6.6 HIGH AVAILABILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Operating System Maturity	Preferred released and QA Tested for >= 6 Months
	Hardware Maturity	Available for purchase for >= 1 Year
	Quality Assurance Certification	Employs a rigorous testing program in a laboratory environment that simulates a "real world" enterprise network. Testing of both hardware and software that verifies feature functionality and interoperability under "real world" traffic loads.

EXPLANATION OF STANDARD

To support the high availability environment it is essential that the maturity of both the hardware and the supporting software be subjected to a rigorous evaluation of a quality assurance program to mitigate the risk of un-scheduled system down time. The solution should be operating in a non-critical environment for a time sufficient enough to identify any flaws or vulnerabilities that may compromise the stability or security of the system. The more time the system has matured will be directly related to the probability that anomalies will be identified and eliminated and there is a corollary to reduced risk of an outage.

EVALUATION FACTORS

- None

IMPLEMENTATION GUIDANCE

- None

3.7 LOAD BALANCER

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Rack Mountable	Required
2	Load Balancing Predictors	Least loaded, Least Bandwidth, Least Connections, Round-Robin, Hash Address
3	Server Health Monitoring	ICMP, TCP, ECHO(tcp&udp), HTTP, HTTPs, FTP, Telnet, DNS, RADIUS, SNMP, Scripted
4	Redundancy	Required
5	Role-Based Administration	Required
6	Throughput	>= 4 Gbps
7	Packets per Second	>= 6.5 million
8	NAT Entries	>= 1 million
9	Connections per Second	>= 325,000 sustained
10	Concurrent Connections	>= 4 million

EXPLANATION OF STANDARD

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

The solution will employ a system that will probe each of the database nodes in both metro data centers to determine which of the metro pairs the target database in currently operating in. The solution will then probe a pool of applications nodes to determine which of the pool within the data center hosting the database are receiving connections for the target data base. The solution will then measure the current loading on each of the

eligible application nodes to determine the least loaded node. The solution will then send connections to the least loaded node.

4 TAXONOMY OF STANDARDS

The primary and secondary attributes identified below are common to all standards sets that will be contained in this standard.

ID	*Primary Attribute*	*Secondary Attribute*
1	Chassis	Type (Modular vs. Static)
		Rack Mounting
		System Throughput
		Packet Processor Redundancy Support
		Control Plane Redundancy Support
		Cooling Redundancy Support
		Power Supply Redundancy Support
		Front-Back Airflow Support
		Online Insertion/Removable Hot-Pluggable Support
2	Packet Processor	Type (Upgradable)
		Layer 3 Switching Throughput (pps)
		Layer 2 Switching Throughput (pps)
		Online Insertion/Removable Hot-Pluggable Support
		Memory
		Protocols switched in hardware
3	Linecards	Interface Speed
		Interface Media Flexibility Support
		ASIC per port (Oversubscription)
		Throughput to Backplane

		Queuing Properties
		Online Insertion/Removable Hot-Pluggable Support
		On-Card distributed packet forwarding support
4	Operating System	Type (Modular)
		Features Supported
		Multi-chassis Link Aggregation
		Multi-chassis Management Interface Aggregation
		Protocols Supported
		Network Management
5	Supportability	Technical Support
		Parts replacement
6	High Availability	Operating System Maturity
		Hardware Maturity
		Quality Assurance Certification
7	Load Balancer	Rack Mountable
		Load Balancing Predictors
		Server Health Monitoring
		Connection Persistence
		Redundancy
		Role-Based Administration
		Throughput
		Packets per Second
		Probes
		NAT Entries
		Connections per Second

		Concurrent Connections
8	Metropolitan Data Center Interconnect (MAN)	Latency
		Transport Method
		Reliability/SLA/Packet Loss
		Path Diversity (Includes Local Loop)
		Carrier Diversity
		Capacity
		Quality of Service
		Networx Carriers
		Traffic Analysis / Probes

APPENDIX A – DEFINITIONS

Click here to enter definitions.

APPENDIX B – REFERENCES

Click here to enter a list of references.

APPENDIX C – ACRONYMS

Refer to the VA Acronym Lookup Web page for a list of VA specific acronyms.

EIE	VA OI&T Enterprise Infrastructure Engineering
MEC	Multi-Chassis Etherchannel
vPC	Virtual Port Channel
VSS	Virtual Switching System
FCoE	Fiber Channel over Ethernet

APPENDIX D – CONTRIBUTORS

The following subject matter experts have contributed to the development of this document as indicated

Name	Organization	Role
Marty Trey	Enterprise Infrastructure Engineering	Team Lead
Tim Dahl	Enterprise Infrastructure Engineering	SME
Gary Benson	Field Operations & Development, Region 1	SME
Mike Cunagin	Contractor representing Region 1	SME
Allen Stevens	Field Operations & Development, Region 1	SME
Hal Haislip	Field Operations & Development, Region 2	SME
Rich Colby	Field Operations & Development, Region 3	SME
Juan Adames	Field Operations & Development, Region 3	SME
Gary Branch	Field Operations & Development, Region 4	SME
Anthony Paul	Field Operations & Development, Region 5	SME
David Polanco	Field Operations & Development, VACO	SME
Casey O'Brien	Field Operations & Development, CDCO	SME
John Mehl	Field Operations & Development, PMO	Stake Holder Representative

John Hina	Field Implementation Services	Stake Holder Representative
Robert Moser	Field Implementation Services	Stake Holder Representative