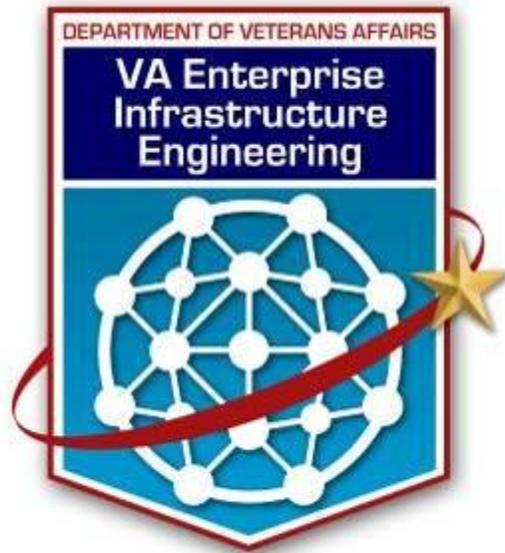




DEPARTMENT OF VETERANS AFFAIRS



OFFICE OF INFORMATION AND TECHNOLOGY
ENTERPRISE INFRASTRUCTURE ENGINEERING

VA Enterprise IT Infrastructure Standard

Wide Area Network

Production 1.0

November 30, 2009

TABLE OF CONTENTS

1	Introduction	1
1.1	Purpose.....	1
1.2	Objectives.....	1
1.3	Scope	1
2	Standards	2
2.1	Class A Router.....	2
2.2	Class B Router.....	4
2.3	Wide Area Network	6
2.3.1	Metro Data Center To Backbone.....	6
2.3.2	Metro Data Center To Region WAn.....	6
2.3.3	Medical Center To Region WAN	7
2.3.4	Medical Center Branch Office To Region WAN	8
2.4	WAN Encryption	9
2.4.1	Data Center	9
2.4.2	Medical Center	9
2.5	WAN Acceleration	10
2.5.1	Data Center	10
2.5.2	Medical Center	11
3	Supporting Details for Standards.....	12
3.1	Class A Router – OC3(qty:4) and higher	12
3.1.1	CHASSIS.....	12
3.1.2	Packet Processor	13
3.1.3	Interface Cards.....	15
3.1.4	Operating System	16

3.1.5	WAN Encryption	18
3.1.6	Platform Supportability.....	20
3.1.7	Platform Maturity.....	21
3.2	Class B Router – DS3(qty:4) up to OC3(qty:3)	22
3.2.1	Chassis	22
3.2.2	Packet Processor	24
3.2.3	Interface Cards	25
3.2.4	Operating System	26
3.2.5	WAN Encryption	28
3.2.6	platform supportability.....	30
3.2.7	Platform maturity.....	31
3.3	Wide Area Network	31
3.3.1	Metro Data Center To Backbone.....	31
3.3.2	Metro Data Center to Region WAN	34
3.3.3	Medical Center to Region WAN.....	35
3.3.4	Medical Center Branch Office to Region WAN.....	39
3.4	WAN Encryption	40
3.4.1	Data Center	40
3.4.2	Medical Center.....	42
3.5	WAN Acceleration	44
3.5.1	Data Center	44
3.5.2	Medical Center.....	47
4	Taxonomy of Standards	50
	Appendix A – Definitions	53
	Appendix B – References	53
	Appendix C – Acronyms	53
	Appendix D – Contributors	53

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 Wide 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
- Provide vendor neutral, or vendor specific where justified, specifications needed to support the design of the VA WAN Infrastructure.

1.3 SCOPE

This standard applies to:

- WAN Infrastructure providing connectivity to the National Data Centers

2 STANDARDS

2.1 CLASS A ROUTER

Class "A" or Carrier class router transporting high volumes of traffic where service disruption is unacceptable.

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type (Modular vs. Static)	Required: Modular
		Throughput	Backplane performance of ≥ 15 Millions packets per second
		Packet Processor Redundancy	Required
		Control Plane Redundancy	Required
		Cooling Redundancy	Required: May be internal to power supplies
		Power Supply Redundancy	Required
		Front-Back Airflow	Required
		Online Insertion & Removal	Required
2	Packet Processor	Type (upgradable)	Required (not integral to chassis)
		Packet Processor Redundancy	Required
		Layer 3 Routing Throughput (PPS)	≥ 15 Million packets per second
		Online Insertion & Removal	Required
		Memory (DRAM, Flash)	Upgradeable - Volatile ≥ 2 GB, Non-Volatile ≥ 1 GB
3	Interface Cards	Interface Speed	Up to 10Gbps
		Interface Media Flexibility	Required (from T-1 to 10Gbps interfaces) T-1, channelized T-1, DS-3, channelized DS-3, OC-3 to OC12, OC-3 to OC-12 Packet over SONET, 10Mbps to 10Gbps Ethernet, at a minimum. (see Evaluation Factors)

ID	Primary Attribute	Secondary Attribute	Specification
		Queuing Properties	Required: see 3.1.3 Operating System
		Online Insertion & Removal	Required
4	Operating System	Type (Modular)	Preferred (See Implementation Guidance)
		Features Supported	Required: Routing, Bridging, Switching
		Routing Protocols	Required: BGPv4, OSPF, Preferred: EIGRP (See Evaluation Criteria)
		Quality of Service	Required: including shaping and policing
		Protocols Supported	Required: MPLS, IPv4, IPv6, 802.1Q, GRE, IPSEC, GDOI
		Network Management	Required: SNMPv1 thru 3, SSH, SSL, IP-SLA, Syslog, Traps (See evaluation criteria for IP Flow export)
5	WAN Encryption	FIPS 140-2 Certification	Required: Complete, or in process. (See Implementation Guidance)
		Throughput	Required: >= 5Gbps (Inbound/Outbound Aggregate) (See Evaluation Criteria)
		Hardware vs Software	Required: Hardware
		Latency	Required: <=100usec
		Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
		Supported Connection Count	Required: >=2000 encrypted tunnels
		Interface Type	Required: Encryption Supported on all interface types
6	Platform Supportability	Technical Support	Required: Phone support 7x24x365
		Parts Replacement	Required: With onsite 4 hour replacement
7	Platform Maturity	Operating System Maturity	Required, 1 year in general production

ID	Primary Attribute	Secondary Attribute	Specification
		Hardware Maturity	Required, 1 year in general production
		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 CLASS B ROUTER

Class “B” Router transports high volumes of traffic, needs redundancy, but does not require higher speed interface density.

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type (Modular vs. Static)	Modular required
		System Throughput	2 Million pps
		Packet Processor Redundancy	Preferred
		Control Plane Redundancy	Preferred
		Cooling Redundancy	Preferred
		Power Supply Redundancy	Required
		Front-Back Airflow	Preferred
		Online Insertion & Removal	Required
2	Packet Processor	Type (upgradable)	Required
		Packet Processor Redundancy	Preferred
		Layer 3 Routing Throughput (pps)	2 Million pps
		Online Insertion & Removal	Preferred
		Memory (DRAM, Flash)	1 GB – DRAM, 256MB - Flash
3	Interface Cards	Interface Speed	Up to multiple OC-3
		Interface Media Flexibility	Modular router must support many interface types

ID	Primary Attribute	Secondary Attribute	Specification
		Queuing Properties	Multiple Queuing strategies
		Online Insertion & Removal	Required
4	Operating System	Type (Modular)	Preferred (See Implementation Guidance)
		Features Supported	Required: Routing, Bridging, Switching
		Routing Protocols	Required: BGPv4, OSPF, Preferred: EIGRP (See Evaluation Criteria)
		Quality of Service	Required: including shaping and policing
		Protocols Supported	Required: MPLS, IPv4, IPv6, 802.1Q, GRE, IPSEC, GDOI
		Network Management	Required: SNMPv1 thru 3, SSH, SSL, IP-SLA, Syslog, Traps (See evaluation criteria for IP Flow export)
5	WAN Encryption	FIPS 140-2 Certification	Required: Complete, or in process. (See Implementation Guidance)
		Throughput	Required: >= 5Gbps (Inbound/Outbound Aggregate) (See Evaluation Criteria)
		Hardware vs Software	Required: Hardware
		Latency	Required: <=100usec
		Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
		Supported Connection Count	Required: >=2000 encrypted tunnels
		Interface Type	Required: Encryption Supported on all interface types
6	Platform Supportability	Technical Support	required
		Parts Replacement	required
7	Platform Maturity	Operating System Maturity	required

ID	Primary Attribute	Secondary Attribute	Specification
		Hardware Maturity	required
		Quality Assurance Certification	required

2.3 WIDE AREA NETWORK

2.3.1 METRO DATA CENTER TO BACKBONE

Specifications identifying National Data Center to VA Backbone connectivity.

ID	Primary Attribute	Secondary Attribute	Specification
1	WAN	Latency	Currently < 60ms
		Transport Method	MPLS Preferred
		Reliability/SLA/Packet Loss	(See Appendix E)
		Path Diversity (Includes Local Loop)	Required
		Carrier Diversity	Required: Qwest and ATT are the carriers of choice based on the Networx Contract
		Capacity	Minimum OC-3 per carrier
		Quality of Service	TBD
		Networx Carriers	Qwest and ATT
		Traffic Analysis / Probes	Analyzed monthly by Capacity Planning Team at the NSOC

2.3.2 METRO DATA CENTER TO REGION WAN

Specifications identifying National Data Center to Region WAN connectivity.

ID	Primary Attribute	Secondary Attribute	Specification
1	WAN	Latency	Require: <= 100ms, Acceptable: <= 60ms, Prefer: <= 40ms

ID	Primary Attribute	Secondary Attribute	Specification
		Transport Method	MPLS Required
		Reliability/SLA/Packet Loss	(See Appendix E)
		Path Diversity (Includes Local Loop)	Required; including diverse LECs, and COs; or CPA as a last resort
		Carrier Diversity	Required; including diverse LECs, and COs; or CPA as a last resort
		Capacity	N+1 (dual-circuits don't exceed 50% utilization)
		Quality of Service	Required, (Priority Queuing, + >= 2 non-priority queues, + default)
		Network Carriers	Qwest and ATT
		Traffic Analysis / Probes	Analyzed monthly by Capacity Planning Team at the NSOC

2.3.3 MEDICAL CENTER TO REGION WAN

Specifications identifying Medical Center to Region WAN connectivity.

ID	Primary Attribute	Secondary Attribute	Specification
1	WAN	Latency	Require: <= 100ms, Acceptable: <= 60ms, Prefer: <= 40ms
		Transport Method	Option #1: MPLS/PTP Hybrid, (See Implementation Guidance) <REGION 2&3> Option #2: MPLS Required, (See Implementation Guidance) <REGION 1&4>
		Reliability/SLA/Packet Loss	(See Appendix E)
		Path Diversity (Includes Local Loop)	Required (direct patient care/continued business operations); Preferred (non patient care) (See Explanation of Standard)
		Carrier Diversity	Required when circuits are MPLS.
		Capacity	N+1 (dual-circuits don't exceed 50% utilization)
		Quality of Service	Required, (Priority Queuing, + >= 2 non-priority queues, +

<i>ID</i>	<i>Primary Attribute</i>	<i>Secondary Attribute</i>	<i>Specification</i>
			default)
		Network Carriers	Required
		Traffic Analysis / Probes	Required

2.3.4 MEDICAL CENTER BRANCH OFFICE TO REGION WAN

Specifications identifying Medical Center Branch Office (OPC, CBOC, Admin Offices) connectivity to Region WAN.

<i>ID</i>	<i>Primary Attribute</i>	<i>Secondary Attribute</i>	<i>Specification</i>
1	WAN	Latency	Require: <= 100ms, Acceptable: <= 60ms, Prefer: <= 40ms
		Transport Method	<u>See Implementation Guidance:</u> Option #1: Direct Pt-to-Pt Option #2: Metro Ethernet Option #3: MPLS
		Reliability/SLA/Packet Loss	(See Appendix E)
		Path Diversity (Includes Local Loop)	Required for Critical Facilities (direct patient care/continued business operations), otherwise Path diversity is preferred
		Carrier Diversity	Not Required
		Capacity	N+1 Required for Critical Facilities, otherwise preferred and determined by business needs
		Quality of Service	Required, (Priority Queuing, + >= 2 non-priority queues, + default)
		Network Carriers	Required
		Traffic Analysis / Probes	Required

2.4 WAN ENCRYPTION

Identifies WAN Encryption requirements for Data Center and Medical Center.

2.4.1 DATA CENTER

ID	Primary Attribute	Secondary Attribute	Specification
1	Encryption	FIPS 140-2 Validation	Complete, or in process. (See Implementation Guidance)
		Throughput	>= 10 Gbps Full Duplex (See Evaluation Criteria)
		Hardware vs. Software	Hardware
		Latency Introduced	<= 100µs
		Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
		Supported Connection Count	>= 10,000
		Cryptography	AES-256 bit
		Jumbo Frame Support	Required
		Interface Type	Carrier Metro Ethernet, Dark-Fiber, and Ethernet over Sonet
		Bypass Capability	Required
2	Chassis	Power Supply Redundancy	Required
		In-Band Management	Required
		Out of Band Management	Recommended

2.4.2 MEDICAL CENTER

ID	Primary Attribute	Secondary Attribute	Specification
-----------	--------------------------	----------------------------	----------------------

ID	Primary Attribute	Secondary Attribute	Specification
1	Encryption	FIPS 140-2 Validation	Complete, or in process. (See Implementation Guidance)
		Throughput	>= 5 Gbps Full Duplex (See Evaluation Criteria)
		Hardware vs. Software	Hardware
		Latency Introduced	<= 100µs
		Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
		Supported Connection Count	>= 2,000 encrypted tunnels
		Cryptography	AES-256 bit, 3DES, IPSec
		Jumbo Frame Support	Required
		Interface Type	Required: Encryption supported on all interface types
		Bypass Capability	Required
2	Chassis	Power Supply Redundancy	Required
		In-Band Management	Required
		Out of Band Management	Recommended

2.5 WAN ACCELERATION

Specifications for WAN Acceleration in the Data Center.

2.5.1 DATA CENTER

ID	Primary Attribute	Secondary Attribute	Specification
1	WAN Acceleration	Interface Type	Copper or Fiber (RJ45 or LC)
		Capacity	See Implementation Guidance
		Implementation (inline vs.	

ID	Primary Attribute	Secondary Attribute	Specification
		redirection)	
		Redundancy	Preferred
		Traffic Types Supported	CIFS, MAPI, MAPI NSPI, HTTP, HTTPS, FTP, MS-SQL, Oracle SQL
		Centralized Management and Reporting	Required
		Connection Management / Visibility	Network monitoring including support for AAA and SNMP is required
		Effectiveness (Application Acceleration, compression, Auto-Discovery)	Greater than 80% data reduction is desired

2.5.2 MEDICAL CENTER

Specifications for WAN Acceleration in the Medical Center

ID	Primary Attribute	Secondary Attribute	Specification
1	WAN Acceleration	Interface Type	Copper or Fiber (RJ45 or LC)
		Capacity	See Implementation Guidance
		Implementation (inline vs. redirection)	
		Redundancy	Preferred
		Traffic Types Supported	CIFS, MAPI, MAPI NSPI, HTTP, HTTPS, FTP, MS-SQL, Oracle SQL
		Centralized Management and Reporting	Required
		Connection Management / Visibility	Network monitoring including support for AAA and SNMP is required
		Effectiveness (Application Acceleration, compression, Auto-Discovery)	Greater than 80% data reduction is desired

3 SUPPORTING DETAILS FOR STANDARDS

3.1 CLASS A ROUTER – OC3(QTY:4) AND HIGHER

3.1.1 CHASSIS

CHASSIS STANDARD

ID	Primary Attribute	Secondary Attribute	Specification
1	Chassis	Type	Required: Modular
		Throughput	Backplane performance of ≥ 15 Millions packets per second
		Packet Processor Redundancy	Required
		Control Plane Redundancy	Required
		Cooling Redundancy	Required: May be internal to power supplies
		Power Supply Redundancy	Required
		Front-Back Airflow	Required
		Online Insertion & Removal	Required

EXPLANATION OF STANDARD

The class “A” router is for use in situations where there is a very high volume of traffic and where any service interruption is unacceptable. It is envisioned that this router would be installed at Data Centers and major network aggregation points within the VA and considered to be at the Core Layer. This type is often referred to as a carrier class router.

A modular chassis will be required in order to support redundant power as well as redundant route processors for required High Availability.

Due to the lack of data to support actual throughput required we have utilized the throughput capabilities of current architectures being utilized today.

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
- WAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

None

3.1.2 PACKET PROCESSOR

STANDARD

<i>ID</i>	<i>Secondary Attribute</i>	<i>Specification</i>
2	Type (upgradable)	Required (not integral to chassis)
	Packet Processor Redundancy	Required
	Layer 3 Routing Throughput (PPS)	>= 15 Million packets per second

	Online Insertion & Removal	Required
	Memory (DRAM, Flash)	Upgradeable - Volatile >= 2GB, Non-Volatile >= 1GB

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 Routing 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 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.

EVALUATION FACTORS

- For throughput and memory attributes, the higher the number, the better.

IMPLEMENTATION GUIDANCE

None

3.1.3 INTERFACE CARDS

STANDARD

ID	Secondary Attribute	Specification
3	Interface Speed	Up to 10Gbps
	Interface Media Flexibility	Required (from T-1 to 10Gbps interfaces) T-1, channelized T-1, DS-3, channelized DS-3, OC-3 to OC12, OC-3 to OC-12 Packet over SONET, 10Mbps to 10Gbps Ethernet, at a minimum. (see Evaluation Factors)
	Queuing Properties	Required: see 3.1.4 Operating System
	Online Insertion & Removal	Required

EXPLANATION OF STANDARD

Interface speeds up to 10 Gbps will need to be supported in a Class A (Carrier Class) Router. The Class A router will reside primarily in the National Data Centers and locations where high volumes of traffic will be routed within the VA.

Interface Media flexibility is important. As requirements change or Carrier Services change it is important to be able to adapt. Today the VA utilizes the above specified interfaces and are required to be supported by the Class A routing platform.

Interface cards are required to support queuing mechanisms instituted by the platforms operating system. See section 3.1.4 for further discussion on queuing.

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

- For Interface Media Flexibility – in order to cover all eventualities a very wide range of interfaces is desired.

IMPLEMENTATION GUIDANCE

None

3.1.4 OPERATING SYSTEM

STANDARD

ID	Secondary Attribute	Specification
4	Type (Modular)	Preferred (See Implementation Guidance)
	Features Supported	Required: Routing, Bridging, Switching
	Routing Protocols	Required: BGPv4, OSPF, Preferred: EIGRP (See Evaluation Criteria)
	Queuing/Quality of Service	Required: including shaping and policing
	Protocols Supported	Required: MPLS, IPv4, IPv6, 802.1Q, GRE, IPSEC, GDOI
	Network Management	Required: SNMPv2c and v3, SSH, SSL, IP-SLA, Syslog, SNMP Traps, Centralized AAA (see evaluation criteria for IP Flow export)

EXPLANATION OF STANDARD

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

The routing platform will be required to support features to include Routing, Bridging, and Switching for flexibility in design for HA environments.

The two routing protocols used within the Department of Veterans Affairs are EIGRP and BGPv4. EIGRP is highly preferred in order to take advantage of existing VA expertise and to be compatible with current infrastructure deployments. OSPF is required as a standards based alternative to EIGRP for compatibility between vendors.

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.

In managing the WAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

- Routing Protocols - The two routing protocols used within the Department of Veterans Affairs are EIGRP and BGPv4. EIGRP is highly preferred in order to take advantage of existing VA expertise and to be compatible with current infrastructure deployments.
- A method to export IP flow information via UDP or SCTP is highly desirable. Since the current VA infrastructure is Cisco today the Netflow feature is utilized to export IP Flow information. The IETF standard is IPFIX which is based on Cisco's Netflow Version 9 implementation.

IMPLEMENTATION GUIDANCE

Modular Operating System – A Modular Operating system that permits the upgrade of code in all modules without service interruption is the goal. Code maturity will be a major consideration on whether or not it is selected. Caveats and known issues identified in manuals, customer feedback, and bug notices will be examined.

3.1.5 WAN ENCRYPTION

STANDARD

ID	Secondary Attribute	Specification
5	FIPS 140-2 Certification	Required: Complete, or in process. (See Implementation Guidance)
	Throughput	Required: >= 5Gbps (Inbound/Outbound Aggregate) (See Evaluation Criteria)
	Hardware vs Software	Required: Hardware
	Latency	Required: <=100usec
	Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
	Supported Connection Count	Required: >=2000 encrypted tunnels
	Interface Type	Required: Encryption Supported on all interface types

EXPLANATION OF STANDARD

When WAN Encryption is required for WAN circuits the encryption method must be FIPS 140-2 certified (see implementation guidance for further details).

Without current data requirements to identify the required encryption throughput the throughput has been identified using the current specifications from technology implemented in the VA today.

Encryption performed in hardware will be a requirement for high volume environments such as Data Centers and VAMC's. Encryption done in software could be acceptable for low bandwidth requirements (see implementation guidance for further details).

Latency requirements were derived from existing infrastructure components deployed in the VA today. There is not enough data to support encryption latency requirements.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead.

EVALUATION FACTORS

- WAN Encryption Throughput – The higher the throughput the better.

IMPLEMENTATION GUIDANCE

FIPS 140-2 Certification - The system must be either FIPS 140-2 NIST certified or at least in stage 3 testing.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead.

3.1.6 PLATFORM SUPPORTABILITY

STANDARD

ID	Secondary Attribute	Specification
5	Technical Support	Required: Phone support 7x24x365
	Parts Replacement	Required: With onsite 4 hour replacement

EXPLANATION OF STANDARD

Due to the HA requirement for a Class A router Technical Support needs to be 7x24x365 with four hour parts replacement.

EVALUATION FACTORS

- If something better than onsite 4 hour support is offered it will be considered.

IMPLEMENTATION GUIDANCE

None

3.1.7 PLATFORM MATURITY

STANDARD

ID	Secondary Attribute	Specification
6	Operating System Maturity	Required, 1 year in general production
	Hardware Maturity	Required, 1 year in general production
	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

Due to the critical nature of a class “A” router, it is imperative for the platform to have been in general production for a period of time in order for the platform to mature and become highly stable. Early platform adoption many times requires frequent OS upgrades due to bugs being found in hardware as well as software.

EVALUATION FACTORS

- The 1 year general production requirement allows past performance to be evaluated.
- Quality Assurance Certification programs employed by vendors.

IMPLEMENTATION GUIDANCE

None

3.2 CLASS B ROUTER – DS3(QTY:4) UP TO OC3(QTY:3)

3.2.1 CHASSIS

STANDARD

ID	Secondary Attribute	Specification
1	Type (Modular vs. Static)	Modular required
	System Throughput	2 Million pps
	Packet Processor Redundancy	Preferred
	Control Plane Redundancy	Preferred
	Cooling Redundancy	Preferred
	Power Supply Redundancy	Required
	Front-Back Airflow	Preferred
	Online Insertion & Removal	Required

EXPLANATION OF STANDARD

A modular chassis will be required in order to support redundant power as well as redundant route processors for required High Availability.

Due to the lack of data to support actual throughput required we have utilized the throughput capabilities of current architectures being utilized today.

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
- WAN Protocols and Features supported
- Experience VA currently has with solution

IMPLEMENTATION GUIDANCE

3.2.2 PACKET PROCESSOR

STANDARD

ID	Secondary Attribute	Specification
2	Type (upgradable)	Required
	Packet Processor Redundancy	Preferred
	Layer 3 Routing Throughput (pps)	2 Million pps
	Online Insertion & Removal	Preferred
	Memory (DRAM, Flash)	1 GB – DRAM, 256MB - Flash

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 Routing 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 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.

EVALUATION FACTORS

- For throughput and memory attributes, the higher the number, the better.

IMPLEMENTATION GUIDANCE

3.2.3 INTERFACE CARDS

STANDARD

ID	Secondary Attribute	Specification
3	Interface Speed	Up to multiple OC-3
	Interface Media Flexibility	Required (from T-1 to OC-3 interfaces) T-1, channelized T-1, DS-3, channelized DS-3, OC-3, OC-3 Packet over SONET, 10 Mbps to 1 Gbps Ethernet, at a minimum. (see Evaluation Factors)
	Queuing Properties	Required: see 3.1.4 Operating System
	Online Insertion & Removal	Required

EXPLANATION OF STANDARD

Interface speeds up to OC-3 will need to be supported in a Class B Router. The Class B router will reside primarily at the Distribution Layer and locations where high volumes of traffic will be routed within the VA.

Interface Media flexibility is important. As requirements change or Carrier Services change it is important to be able to adapt. Today the VA utilizes the above specified interfaces and are required to be supported by the Class A routing platform.

Interface cards are required to support queuing mechanisms instituted by the platforms operating system. See section 3.2.4 for further discussion on queuing.

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

- For Interface Media Flexibility – in order to cover all eventualities a very wide range of interfaces is desired.

IMPLEMENTATION GUIDANCE

3.2.4 OPERATING SYSTEM

STANDARD

ID	Secondary Attribute	Specification
4	Type (Modular)	Preferred (See Implementation Guidance)
	Features Supported	Required: Routing, Bridging, Switching
	Routing Protocols	Required: BGPv4, OSPF, Preferred: EIGRP (See Evaluation Criteria)
	Quality of Service	Required: including shaping and policing
	Protocols Supported	Required: MPLS, IPv4, IPv6, 802.1Q, GRE, IPSEC, GDOI
	Network Management	Required: SNMPv2c and v3, SSH, SSL, IP-SLA, Syslog, SNMP Traps, Centralized AAA (see evaluation criteria for IP Flow export)

EXPLANATION OF STANDARD

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

The routing platform will be required to support features to include Routing, Bridging, and Switching for flexibility in design for HA environments.

The two routing protocols used within the Department of Veterans Affairs are EIGRP and BGPv4. EIGRP is highly preferred in order to take advantage of existing VA expertise and to be compatible with current infrastructure deployments. OSPF is required as a standards based alternative to EIGRP for compatibility between vendors.

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.

In managing the WAN infrastructure the specified protocols are required to support roles based management, configuration management, and monitoring.

EVALUATION FACTORS

- Routing Protocols - The two routing protocols used within the Department of Veterans Affairs are EIGRP and BGPv4. EIGRP is highly preferred in order to take advantage of existing VA expertise and to be compatible with current infrastructure deployments.
- A method to export IP flow information via UDP or SCTP is highly desirable. Since the current VA infrastructure is Cisco today the Netflow feature is utilized to export IP Flow information. The IETF standard is IPFIX which is based on Cisco's Netflow Version 9 implementation

IMPLEMENTATION GUIDANCE

Modular Operating System – A Modular Operating system that permits the upgrade of code in all modules without service interruption is the goal. Code maturity will be a major consideration on whether or not it is selected. Caveats and known issues identified in manuals, customer feedback, and bug notices will be examined.

3.2.5 WAN ENCRYPTION

STANDARD

ID	Secondary Attribute	Specification
5	FIPS 140-2 Certification	Required: Complete, or in process. (See Implementation Guidance)
	Throughput	Required: >= 900Mbps (Inbound/Outbound Aggregate) (See Evaluation Criteria)
	Hardware vs Software	Required: Hardware
	Latency	Required: <=100usec
	Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
	Supported Connection Count	Required: >=2000 encrypted tunnels
	Interface Type	Required: Encryption Supported on all interface types

EXPLANATION OF STANDARD

When WAN Encryption is required for WAN circuits the encryption method must be FIPS 140-2 certified (see implementation guidance for further details).

Without current data requirements to identify the required encryption throughput the throughput has been identified using the current specifications from technology implemented in the VA today.

Encryption performed in hardware will be a requirement for high volume environments such as Data Centers and VAMC's. Encryption done in software could be acceptable for low bandwidth requirements (see implementation guidance for further details).

Latency requirements were derived from existing infrastructure components deployed in the VA today. There is not enough data to support encryption latency requirements.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead.

EVALUATION FACTORS

- WAN Encryption Throughput – The higher the throughput the better.

IMPLEMENTATION GUIDANCE

FIPS 140-2 Certification - The system must be either FIPS 140-2 NIST certified or at least in stage 3 testing.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead.

3.2.6 PLATFORM SUPPORTABILITY

STANDARD

ID	Secondary Attribute	Specification
5	Technical Support	Required: Phone support 7x24x365
	Parts Replacement	Required: With onsite 4 hour replacement

EXPLANATION OF STANDARD

Due to the HA requirement for a Class B router Technical Support needs to be 7x24x365 with four hour parts replacement.

EVALUATION FACTORS

- If something better than onsite 4 hour support is offered it will be considered.

IMPLEMENTATION GUIDANCE

3.2.7 PLATFORM MATURITY

STANDARD

ID	Secondary Attribute	Specification
6	Operating System Maturity	required
	Hardware Maturity	required
	Quality Assurance Certification	required

EXPLANATION OF STANDARD

Due to the critical nature of a Class “B” router, it is imperative for the platform to have been in general production for a period of time in order for the platform to mature and become highly stable. Early platform adoption many times requires frequent OS upgrades due to bugs being found in hardware as well as software.

EVALUATION FACTORS

- The 1 year general production requirement allows past performance to be evaluated.
- Quality Assurance Certification programs employed by vendors.

IMPLEMENTATION GUIDANCE

3.3 WIDE AREA NETWORK

3.3.1 METRO DATA CENTER TO BACKBONE

STANDARD

ID	Secondary Attribute	Specification
1	Latency	Currently < 60ms
	Transport Method	MPLS Preferred
	Path Diversity (Includes Local Loop)	Required

	Carrier Diversity	Qwest and ATT are the carriers of choice based on the Network Contract
	Capacity	Minimum OC-3 per carrier
	Quality of Service	TBD
	Network Carriers	Qwest and ATT
	Traffic Analysis	Analyzed monthly by Capacity Planning Team at the NSOC

EXPLANATION OF STANDARD

A 2-port OC-3 adaptor will be used to terminate the OC-3 circuits from AT&T and Sprint on the One-VA distribution routers. Each AS will have two exit points out of the Region via a circuit to each MPLS network. Equal cost routing will take place at the distribution nodes and traffic will utilize and load share between both MPLS connections.

Each One-VA distribution router will have an External BGP (eBGP) peer with both MPLS networks (a peer with AT&T AS 7018 and a peer with Qwest AS 209). The routers will also have an internal BGP (iBGP) peer with each other to exchange routing information. The routers will be configured to forward only the networks that belong to that region. This will be accomplished with the use of route maps and prefix lists defining the prefixes belonging to a region. The routes will be advertised from each distribution node to each of the carrier MPLS networks with equal metrics per demarc location. Routes will be advertised back into other VA regions from both MPLS networks.

The routing architecture will make use of route maps to set metrics and to tag routes with specific BGP communities. The current ATM routing topology uses the Multi-Exit Discriminator (MED) attribute to set metrics. The MED attribute tells external neighbors about the preferred path into an AS when there are multiple entry points into the AS. The exit point with the lowest metric is preferred. If a MED is received over an external BGP link, it is propagated over internal links to other BGP speaking routers within the AS. MED is a non-transitive attribute and is not passed from the provider networks; therefore, along with setting MEDs using routes maps, routes will also be tagged with specific communities. The communities are a transitive attribute and will be used to mark or tag specific subnets. All routes inbound from the MPLS networks to a region will have a specific community associated with them. This will allow the routes to be identified and metrics set according to specific route policies. The router configurations will set metrics inbound and outbound for routes using communities. The VA Regions will be required to split their networks so one-half of their routes prefer one distribution node while the remaining routes will prefer the second distribution node.

Each distribution router will have two WAN connections and will receive all routes on each; however, BGP, by default, uses the best path algorithm to install the one best path in the IP routing table to use for traffic forwarding. If BGP multipath is enabled using the *maximum-paths* configuration command and the eBGP paths are learned from the same neighboring AS, instead of picking one best path, multiple paths are installed in the route table. However, the One-VA distribution routers will receive two equal eBGP paths from different neighboring AS numbers (AS 209 for Qwest and AS 7018 for AT&T). BGP will not install the desired equal cost routes based on the fact that the routes are being sent by different AS. In order to overcome this "limitation" and accomplish equal cost routing at each distribution node, the routing configuration will make use of a new BGP command: *bgp best path as-path multipath-relax*. Based on lab testing conducted by the engineering team, this configuration

command works by installing equal paths from differing AS into the Routing Information Base (RIB). This will allow traffic to effectively load share between the MPLS networks. This works with the caveat that the router BGP configuration command *bgp deterministic-med* is not configured.

Default Route and Internet Gateway Tunnel

Each tunnel will be built from the loopback0 address of the distribution router to a unique loopback1 address on one of the gateway routers in each location. The loopback1 addresses will be independent of the loopback0 addresses used for the gateway routers administration. By using a separate loopback address at the gateways recovering the tunnels in the case of a catastrophic failure of the primary tunnel endpoint router hardware can be accomplished by rebuilding the tunnels to the remaining router in service using the loopback1 information from the failed router.

Since the tunnel overlay is already engineered and established it can be used for IPv6 routing. Once the decision to enable IPv6 in a dual protocol stack routing configuration is made, the BGP peers to the gateway locations will be modified to allow the IPv4 default gateway as well as the full IPv6 routing table to pass over the tunnel peers. Once the MPLS carriers can support the IPv6 protocol, the tunnels will be reconfigured to allow only the default route for IPv4 to pass over them. The default IPv6 route will also be added to the tunnel design if all 4 gateway locations can support that configuration.

Fail-over Design

The tunnel and default routing failover design was engineered to provide for failover if the default route for a gateway location is no longer received. If the route is not received by the preferred gateway, the router will no longer announce a default route to the distribution node. Since the Distribution nodes are split between gateway locations, the other regional distribution node will continue to receive the default route and outbound traffic will shift to that distribution node.

In the case of a distribution node's loss of both MPLS providers the tunnel would by default attempt to establish through the other distribution point. Since this is an undesirable failover situation the regions BGP peer between the two distribution nodes will filter out the loopback1 gateway routes it learns from the MPLS carriers. With these routes filtered if a distribution node does lose both carriers the tunnel will fail and the remaining default route will remain as the only default route for the region.

EVALUATION FACTORS

-

IMPLEMENTATION GUIDANCE

3.3.2 METRO DATA CENTER TO REGION WAN

STANDARD

ID	Secondary Attribute	Specification
1	Latency	Require: <= 100ms, Acceptable: <= 60ms, Prefer: <= 40ms
2	Transport Method	MPLS Required
3	Reliability/SLA/Packet Loss	(See Appendix E)
4	Path Diversity (Includes Local Loop)	Required; including diverse LECs, and COs; or CPA as a last resort
5	Carrier Diversity	Required
6	Capacity	N+1 (dual-circuits don't exceed 50% utilization)
7	Quality of Service	Required, (Priority Queuing, + >= 2 non-priority queues, + default)
8	Network Carriers	Required
9	Traffic Analysis / Probes	Required

EXPLANATION OF STANDARD

The Network contract stipulates less than 100ms of latency. It is preferred to achieve less than 60ms of latency on the WAN due to observed performance on existing networks with a maximum round trip of 60ms. The requirement for 40ms is an un-verified assumption of a VistA requirement.

Data Center connectivity to Region WAN will be MPLS assuming Carrier PE node location to the Data Center can support application performance requirements. With an MPLS network the technology lends itself to physical media diversity.

Reliability/SLA/Packet Loss has been pulled from the Network contract and is included below and applies to their MPLS Service offering.

Carrier diversity is required to protect against a catastrophic carrier network outage. Extra attention to ensure path diversity would be required. To ensure physical path diversity in a multi carrier design, the carriers must work together to assure the paths are in fact diverse. Physical path diversity may prove to be difficult to maintain in a multi carrier design due to carrier reroutes during carrier outages. Physical path redundancy includes local loop including diverse Carrier Central Offices for the Local Exchange Carrier and diverse Campus Fiber paths with physical separation of fiber paths of at least 25 feet.

N+1 capacity allows for survivability of a circuit failure without impacting performance.

Quality of Service (QoS) is limited to the capabilities offered by the MPLS providers on the Network Contract. Qwest and AT&T both support four QoS queues.

The GSA Network Contract identifies the Telecommunications Carriers to be utilized. Primary MPLS carrier will be AT&T and the secondary carrier will be Qwest.

Traffic Analysis will be required for management, support, and capacity planning. There are multiple methods in implementing a traffic analysis solution. The traffic analysis solution will be impacted by Data Center design.

EVALUATION FACTORS

- Latency (Lower is better)
- QoS, more queues are better

IMPLEMENTATION GUIDANCE

3.3.3 MEDICAL CENTER TO REGION WAN

STANDARD

<i>ID</i>	<i>Secondary Attribute</i>	<i>Specification</i>
1	Latency	Require: <= 100ms, Acceptable: <= 60ms, Prefer: <= 40ms
2	Transport Method	Option #1: MPLS/PTP Hybrid, (See Implementation Guidance) <REGION 2&3> Option #2: MPLS Required, (See Implementation Guidance) <REGION 1&4>
3	Reliability/SLA/Packet Loss	See Appendix E
4	Path Diversity (Includes Local Loop)	Required; including diverse LECs, and COs; or CPA as a last resort
5	LD Carrier Diversity	Required when circuits are MPLS.
6	Capacity	N+1 (dual-circuits don't exceed 50% utilization)
7	Quality of Service	Required, (Priority Queuing, + >= 2 non-priority queues, + default)
8	Network Carriers	Required
9	Traffic Analysis / Probes	Required

EXPLANATION OF STANDARD

The Networkx contract stipulates less than 100ms of latency. It is preferred to achieve less than 60ms of latency on the WAN due to observed performance on existing networks with a maximum round trip of 60ms. The requirement for 40ms is an un-verified assumption of a VistA requirement.

Where feasible path diversity will be required for Medical Center connectivity to a Region WAN or National Data Center.

Carrier diversity is required to protect against a catastrophic carrier network outage. Extra attention to ensure path diversity would be required. To ensure physical path diversity in a multi carrier design, the carriers must work together to assure the paths are in fact diverse. Physical path diversity may prove to be difficult to maintain in a multi carrier design due to carrier reroutes during carrier outages. Physical path redundancy includes local loop including diverse Carrier Central Offices for the Local Exchange Carrier and diverse Campus Fiber paths with physical separation of fiber paths of at least 25 feet.

Capacity needs to be able to survive a circuit failure without impacting performance.

EVALUATION FACTORS

- Latency (Lower is better)
- QoS, more queues are better

IMPLEMENTATION GUIDANCE

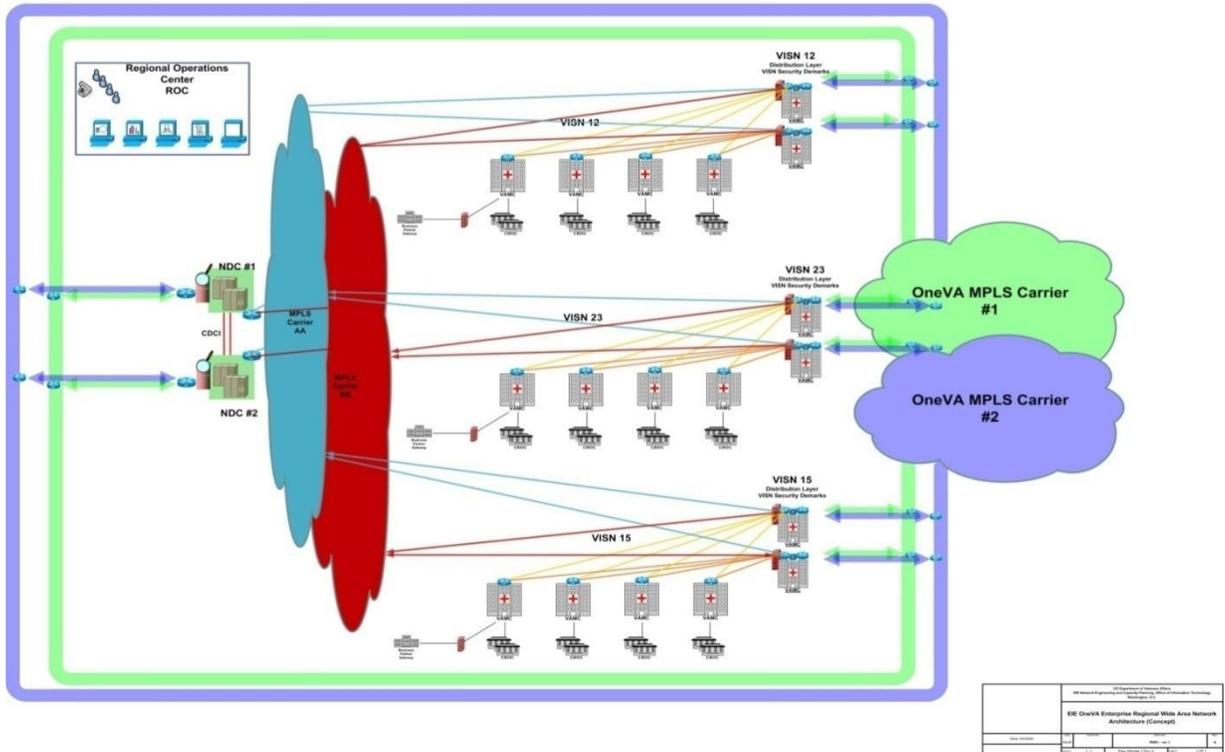
Option #1 for VISNs within Regions 2 and 3 will leverage the performance, additional QoS queues, and security offered by Point-To-Point (PTP) circuits within a VISN's boundaries, and the cost saving associated MPLS over longer distances to the National Data Center (NDC) facilities. Specifically, within the VISN, option #1 calls for one full-line rate PTP DS3 circuit from each of the VISN's non-demark facilities to each demark facilities within the VISN (in most cases this means two DS3 will be installed). The Networkx vendor shall ensure the two PTP DS3s are path-diverse from each non-demark facility to each demark facility; it is assumed where financially prudent that physical diverse entry paths will be built-out where not already available. Industry best practices leveraging Differentiated Services Code Point (DSCP) shall be implemented to ensure a viable Voice, Video, and Data Quality of Service (QoS) solution is provided within each VISN.

Traffic flow between facilities within Region 2 or 3 and the NDC will be facilitated with carrier-redundant MPLS connections at each of the VISN's demark locations. At each VISN demark two routers will be installed, one for each carrier. Each Region 2 and 3 NDC will also have associated carrier-redundant connections to both providers MPLS networks. Specifically, to ensure maximum up-time, each VISN demark facility will have connections to both Networkx provider's MPLS networks, and each metro data center in both regions will also have connections into both providers. For scalability purposes, initial expectations for CE-to-PE connections are OC-3 connections for demark facility-to-provider CE-to-PE connections, and OC-12 for NDC-to-provider CE-to-PE connections.

The diagram provided below is included as an example of this offering.



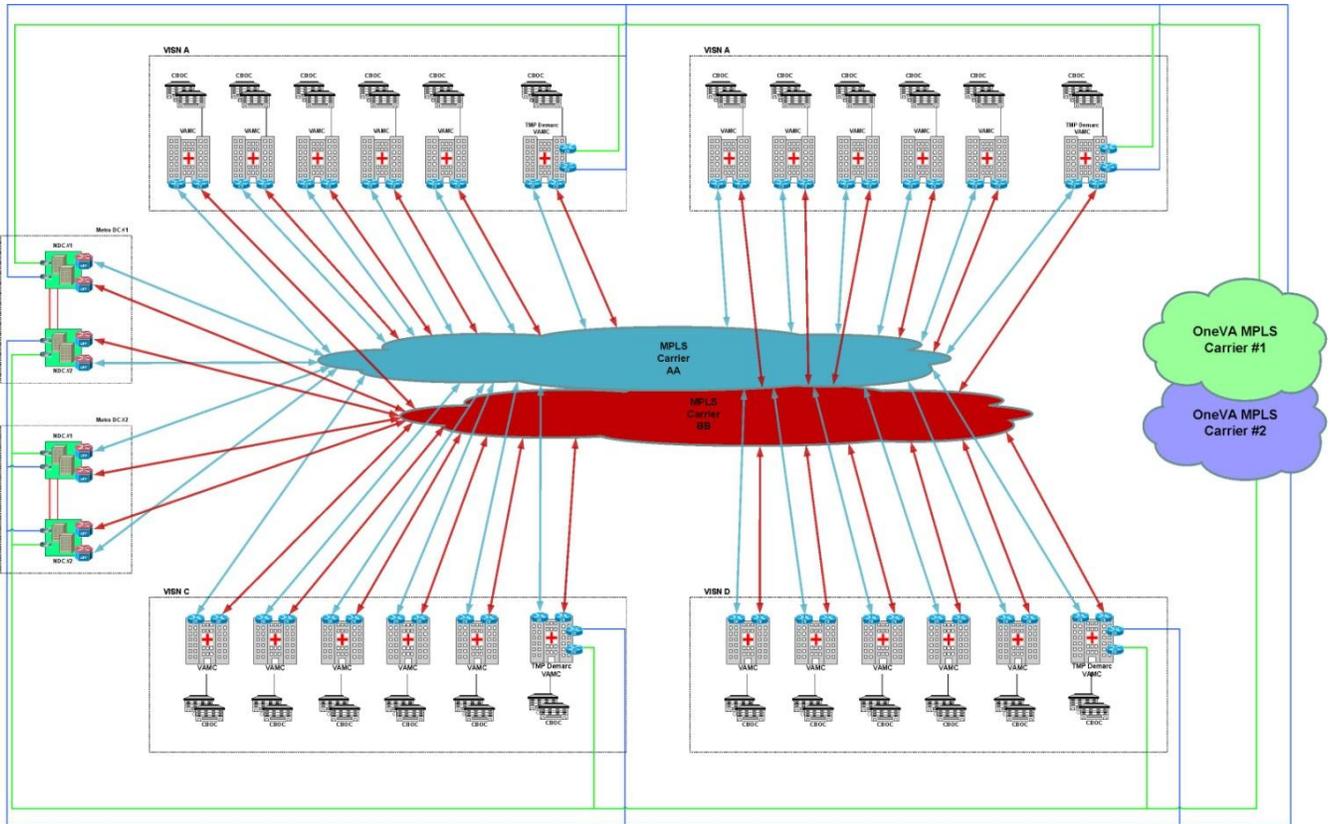
ONE VA Enterprise Regional Wide Area Network Conceptual Architecture
(Partial Region 2 VISNs Represented)



Option #2 for VISNs within Regions 1 and 4 will consist of redundant connections (2) to two Networkx vendor provided MPLS networks (VRF's) from each Medical Center. The total WAN bandwidth capacity to the Medical Center and Datacenter facilities shall be N+1 to provide no degradation in service in the event of one circuit failure. Termination of each MPLS circuit shall be made to redundant router equipment. The Networkx vendor shall ensure the two MPLS circuits are path-diverse from each non-demark facility to each demark facility; it is assumed where financially prudent that physical diverse entry paths will be built-out where not already available. The two Med Center MPLS CE-to-PE circuits will either be DS-3 or OC-3 depending on the bandwidth capacity requirements of the facility. The Datacenter MPLS CE-to-PE circuits will be OC-12.

The intent of this design is to utilize both MPLS carrier networks as redundant equal cost paths for all traffic within each Region. This includes VAMC to Datacenter and VAMC to VAMC. Industry best practices leveraging Differentiated Services Code Point (DSCP) shall be implemented to ensure a viable Voice, Video, and Data Quality of Service (QoS) solution is provided within each VISN/Region.

The diagram provided below is included as an example of this offering.



3.3.4 MEDICAL CENTER BRANCH OFFICE TO REGION WAN

STANDARD

ID	Secondary Attribute	Specification
1	Latency	Require: <= 100ms, Acceptable: <= 60ms, Prefer: <= 40ms
2	Transport Method	<u>See Implementation Guidance:</u> Option #1: Direct Pt-to-Pt Option #2: Metro Ethernet Option #3: MPLS
3	Reliability/SLA/Packet Loss	(See Appendix E)
4	Path Diversity (Includes Local Loop)	Required (direct patient care/continued business operations); Preferred (non patient care) (See Explanation of Standard)
5	Carrier Diversity	Not Required
6	Capacity	See Implementation guidance
7	Quality of Service	Required, (Priority Queuing, + >= 2 non-priority queues, + default)
8	Network Carriers	Required
9	Traffic Analysis / Probes	Required

EXPLANATION OF STANDARD

The Network contract stipulates less than 100ms of latency. It is preferred to achieve less than 60ms of latency on the WAN due to observed performance on existing networks with a maximum round trip of 60ms. The requirement for 40ms is an un-verified assumption of a VistA requirement.

Carrier Services available today and in use by the VA consist of dedicated Point to Point Circuits, Metro Ethernet/Dark Fiber, and MPLS. See implementation guide for further details.

Reliability/SLA/Packet Loss has been pulled from the Network contract and is included below and applies to their MPLS Service offering.

Where feasible path diversity will be required for direct patient care facilities or facilities that have been deemed critical to continued business operations.

Carrier diversity where required protects against a catastrophic carrier network outage. Extra attention to ensure path diversity would be required. To ensure physical path diversity in a multi carrier design, the carriers must work together to assure the paths are in fact diverse. Physical path diversity may prove to be difficult to maintain in a multi carrier design due to carrier reroutes during carrier outages. Physical path redundancy includes local loop including diverse Carrier Central Offices for the Local Exchange Carrier and diverse Campus Fiber paths.

N+1 capacity allows for survivability of a circuit failure without impacting performance. N+1 is required for direct patient care facilities or where business needs dictate reduced service is not acceptable for business operations to continue.

Quality of Service (QoS) is limited to the capabilities offered by the MPLS providers on the Network Contract. Qwest and AT&T both support four QoS queues. QoS capabilities for Metro Ethernet/Dark Fiber and dedicated point to point circuits will only be limited to vendor transport equipment and will need to meet minimum requirements of the MPLS services available.

The GSA Network Contract identifies the Telecommunications Carriers to be utilized. Primary MPLS carrier will be AT&T and the secondary carrier will be Qwest.

Traffic Analysis will be required for management, support, and capacity planning. There are multiple methods in implementing a traffic analysis solution. The traffic analysis solution will be impacted by Data Center design.

Capacity needs to be able to survive a circuit failure without impacting performance

EVALUATION FACTORS

- Latency (Lower is better)
- QoS, more queues are better

IMPLEMENTATION GUIDANCE

Within the Region WAN the transport method will greatly be determined by latency requirements of existing applications deployed, Carrier Services available, and density of Carrier MPLS PE Nodes.

3.4 WAN ENCRYPTION

3.4.1 DATA CENTER

STANDARD

<i>ID</i>	<i>Secondary Attribute</i>	<i>Specification</i>
1	FIPS 140-2 Validation	Required: Complete, or in process. (See Implementation Guidance)
	Throughput	>= 10 Gbps Full Duplex (See Evaluation Criteria)

	Hardware vs. Software	Hardware
	Latency Introduced	<= 100µs
	Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
	Supported Connection Count	>= 10,000
	Cryptography	AES-256 bit
	Jumbo Frame Support	Required
	Interface Type	Carrier Metro Ethernet, Dark-Fiber, and Ethernet over Sonet
	Bypass Capability	Required
2	Power Supply Redundancy	Required
	In-Band Management	Required
	Out of Band Management	Recommended

EXPLANATION OF STANDARD

When WAN Encryption is required between data centers, the encryption method must be FIPS 140-2 certified (see implementation guidance for further details).

Encryption throughput must satisfy high speed LAN interfaces typically found in the datacenter. Because of the expected proliferation of 10GB in the datacenter, any encryption technology procured for the data center must support 10GB at full duplex with little added latency.

At speeds of 10GB, encryption can only be performed in hardware and therefore hardware based encryption is a requirement for high volume environments such as Data Centers and VAMC's.

Latency requirements were derived from existing applications and infrastructure components deployed in the VA today. There is not enough data to support specific encryption latency requirements.

Redundant power supplies are required to satisfy the data center high availability (HA) design requirement. If the encryption device fails it should fail over to bypass mode ensuring no interruption to services or data.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead.

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.

EVALUATION FACTORS

- WAN Encryption Throughput – The higher the throughput the better.
- AES-256 is the preferred Cryptographic algorithm although 3DES is required to meet FIPS 140-2 criteria.
- Separate Hardware Appliances – Separating the Encryption platform from the Routing hardware is preferred. This prevents a router’s IOS upgrade from voiding a product’s FIPS 140-2 certification. FIPS 140-2 Certification - The system must be either FIPS 140-2 NIST certified or at least in stage 3 testing.

IMPLEMENTATION GUIDANCE

FIPS 140-2 Certification - The system must be either FIPS 140-2 NIST certified or at least in stage 3 testing.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead

3.4.2 MEDICAL CENTER

STANDARD

ID	Secondary Attribute	Specification
1	FIPS 140-2 Certification	Required: Complete, or in process. (See Implementation Guidance)
	Throughput	>= 5 Gbps Full Duplex (See Evaluation Criteria)
	Hardware vs Software	Hardware
	Maximum Added Latency	<= 100µs
	Performed using Encapsulation(Tunneling) vs. Payload	Payload preferred, (See Implementation Guidance)
	Supported Connection Count	>= 2,000 encrypted tunnels
	Cryptography	AES-256 bit, 3DES, IPSec
	Bypass Capability	Required
	Jumbo Frame Support	Required

	Interface Type	Encryption Supported on all interface types
2	Power Supply Redundancy	Required
	In-Band Management	Required
	Out of Band Management	Recommended

EXPLANATION OF STANDARD

Encryption hardware used at medical centers must be FIPS 140-2 certified (see implementation guidance for further details).

Encryption hardware must be capable of connecting to all of the different WAN and LAN interface types found at the medical center. Encryption throughput capabilities must satisfy the aggregate throughput of all the different interfaces found at the medical center.

Because of the volume of data processed at the medical centers, hardware encryption is a requirement at VAMC's.

Latency requirements were derived from existing applications and infrastructure components deployed in the VA today. There is not enough data to support specific encryption latency requirements.

If the encryption device fails it should fail over to bypass mode ensuring no interruption to services or data.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead.

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

EVALUATION FACTORS

- WAN Encryption Throughput – The higher the throughput the better.
- AES-256 is the preferred Cryptographic algorithm although 3DES is required to meet FIPS 140-2 criteria.
- Separate Hardware Appliances – Although separating encryption from the routing platform is preferred; leveraging existing router hardware by adding encryption modules will satisfy the agency's encryption requirements. However, integrated routing and encryption does create additional challenges to maintaining FIPS 140-2 certification

IMPLEMENTATION GUIDANCE

FIPS 140-2 Certification - The system must be either FIPS 140-2 NIST certified or at least in stage 3 testing.

Encryption using Encapsulation – Payload encryption is preferred as it preserves the packet headers that include routing and Quality of Service (QoS) parameters. This reduces management and configuration overhead

3.5 WAN ACCELERATION

3.5.1 DATA CENTER

STANDARD

ID	Secondary Attribute	Specification
1	Interface Type	Copper or Fiber (RJ45 or LC)
	Capacity	See Implementation Guidance
	Implementation (inline vs. redirection)	
	Redundancy	Preferred
	Traffic Types Supported	CIFS, MAPI, MAPI NSPI, HTTP, HTTPS, FTP, MS-SQL, Oracle SQL
	Centralized Management and Reporting	Required
	Connection Management / Visibility	Network monitoring including support for AAA and SNMP is required
	Effectiveness (Application Acceleration, compression, Auto-Discovery)	Greater than 80% data reduction is desired

EXPLANATION OF STANDARD

Interface type can be Copper (RJ45) or Fiber (LC/SC/ST)

The system will require having redundant connections to the network in case of an interface failure on either side of the connection.

The system is required to accelerate or optimized the following traffic types/Applications: CIFS, MAPI, MAPI NSPI, HTTP, HTTPS, FTP, MS-SQL, Oracle SQL, (Exchange 2003, 2007 is considered MAPI)

Centralized Management and reporting will be required to help with managing the solution as well as troubleshooting in an efficient manner.

Network monitoring, including support for AAA and SNMP, is required for these systems.

Effectiveness (Application Acceleration, compression, Auto-Discovery) – Greater than 80% data reduction is desired. Also, HTTP caching, separate from byte level caching is desired. . Data compression is required, normally using both layer 6 and 7 servicing of applications. Different compression algorithms can be used including LZ compression. Data Store or Byte level caching storage capability of 3 Terabytes or higher is required. This allows for a FIFO queue technique for byte cache instances to stay in the optimizer for as long as possible in case that cache instance is referenced again, which would keep it off the Wan circuit. Auto discovery allows the appliances to create optimized connections with other WAN optimizers in a dynamic fashion, removing the need to statically set up the peering between med centers and data center throughout the VA.

Redirector device –

Interface type – RJ-45 or LC (Fiber) – Both types need to be supported because both are used when connecting the Wan router to the core switches. RJ-45 are far cheaper because the failure to wire option needed for device failure is cheaper to accomplish for copper than it is for fiber (light).

Interface bypass port count – interface count needs to be able to support 4 Ethernet segments. Each segment between the WAN router and the Core switch will need a WAN and LAN port interface. A minimum of 8 ports are the current requirement.

Implementation – WAN optimization systems will be required to be put in line to optimize the traffic seen on the segments connecting the WAN routers to the LAN core switches. The redirector takes the place of the WAN optimization systems in the data centers, thereby requiring them to be in line.

Interface Capability – Interfaces are required to have the fail to wire or fail open capability in the case of power loss or system failure. Traffic then travels through the redirector without affect.

Link State Propagation -- To ensure consistency throughout networking monitoring systems, redirector appliances will be required to propagate a link state change to an opposing network device.

Centralized Management and Reporting – Required, usually by a management appliance.

Connection Management/ Visibility – Network monitoring including support for AAA and SNMP is required for these systems.

Redirection Peer – The systems redirect traffic to the cluster of WAN optimization systems for servicing so being able to detect whether the peer has failed is required.

Multiple Peer Accelerator/Optimizer Support – the redirector is required to support multiple peers which can be added to or subtracted from during production without detrimental affect.

Load Balancing Support – The redirector is required to be able to balance the load between the peer devices according to multiple algorithms, including load levels, byte caching warmth or availability or in preparations to remove or add peer devices.

Seamless Peer Additions – The redirector is required to support the addition or subtraction of peer devices to increase the load capability or to remove failed devices without detrimental affect to production environment.

Seamless Peer failover – Redirector will support failover from failed devices without negative affect to optimization process.

Ether channel support – Since we do not use this between WAN router and Core switch this is not needed.

EVALUATION FACTORS

- Ease of installation and maintenance is of prime importance in choosing a standard for WAN Optimization. Ease of maintenance is the most important factor as problems can develop with applications and traffic types that may stop users from being able to do their job.
- Visibility of connections statistics and pass through variables can help in the trouble shooting of problems being caused by the WAN optimizers. The manageability of these systems is just as important as the benefits they provide.
- Vendor support needs to be considered as evaluating factor. This can be ascertained by talking to other customers and searching the Internet for testimonial information.
- Though not a requirement, transparency mode is useful when looking at a WAN Optimization standard. This allows the IP addresses of the actual server/client will be preserved instead of the traffic being tunneled using the IP addresses of the Wan optimizers.
- WAN monitoring and reporting as well as QOS schemes can run in the same manner before the installation of WAN Optimizers. Again, this is not required because most optimizers provide a method of doing these things but it should be considered.]
- Due to the environment of the data center, it is imperative to provide a method to effectively maintain the WAN Optimizers in a safe fashion and to provide a tolerant manner to replace failed devices and increase capability without any network interruptions. Redirectors must also provide effective and thorough load balancing between the WAN optimizers. Load levels, different sized machines working together, and data warmth (up to date and easily accessible byte cache instances) must all be taken into consideration.
- Keeping the redirector separate from the routers and switches will help avoid IOS bugs that can and will materialize, will help lower the load on these switches, and will make trouble shooting optimization issues far easier.

IMPLEMENTATION GUIDANCE

Capacity will need to be determined by calculating the required concurrent number of workstation connections to the Data Center that needs acceleration. Without Exchange, recommended estimations have been identified 5 to 10 connections per workstation plus 15% growth as a guide to determine necessary capacity at the Data Center.

Many factors need to be considered while deploying and maintaining WAN optimizers. Many applications evolve or are upgraded on the client and server systems and this can have a detrimental effect on the method and effectiveness of the way the WAN optimizer operates. For instance, when MS Office was upgraded to 2007, encryption was turned on in all client systems in regards to their connections to the Exchange server. Most optimizers cannot effectively accelerate or compress encrypted or previously compressed data going across the wire. The staff that is responsible for the operations of the Optimizers must participate in all groups involved in changing profiles and applications on servers and clients. SMB signing and SMBv2 is another example of data types

that may cause issues. It will be important to use WAN optimizers that can handle these kinds of security requirements and changes.

3.5.2 MEDICAL CENTER

STANDARD

ID	Secondary Attribute	Specification
1	Interface Type	Copper or Fiber (RJ45 or LC)
	Capacity	See Implementation Guidance
	Implementation (inline vs. redirection)	
	Redundancy	Preferred
	Traffic Types Supported	CIFS, MAPI, MAPI NSPI, HTTP, HTTPS, FTP, MS-SQL, Oracle SQL
	Centralized Management and Reporting	Required
	Connection Management / Visibility	Network monitoring including support for AAA and SNMP is required
	Effectiveness (Application Acceleration, compression, Auto-Discovery)	Greater than 80% data reduction is desired

EXPLANATION OF STANDARD

Interface type can be Copper (RJ45) or Fiber (LC/SC/ST)

The system will require having redundant connections to the network in case of an interface failure on either side of the connection.

The system is required to accelerate or optimized the following traffic types/Applications: CIFS, MAPI, MAPI NSPI, HTTP, HTTPS, FTP, MS-SQL, Oracle SQL, (Exchange 2003, 2007 is considered MAPI)

Centralized Management and reporting will be required to help with managing the solution as well as troubleshooting in an efficient manner.

Network monitoring, including support for AAA and SNMP, is required for these systems.

Effectiveness (Application Acceleration, compression, Auto-Discovery) – Greater than 80% data reduction is desired. Also, HTTP caching, separate from byte level caching is desired. . Data compression is required, normally using both layer 6 and 7 servicing of applications. Different compression algorithms can be used including LZ compression. Data Store or Byte level caching storage capability of 3 Terabytes or higher is required. This allows for a FIFO queue technique for byte cache instances to stay in the optimizer for as long as possible in case that cache instance is referenced again, which would keep it off the Wan circuit. Auto discovery allows the appliances to create optimized connections with other WAN optimizers in a dynamic fashion, removing the need to statically set up the peering between med centers and data center throughout the VA.

EVALUATION FACTORS

- Ease of installation and maintenance is of prime importance in choosing a standard for WAN Optimization. Ease of maintenance is the most important factor as problems can develop with applications and traffic types that may stop users from being able to do their job.
- Visibility of connections statistics and pass through variables can help in the trouble shooting of problems being caused by the WAN optimizers. The manageability of these systems is just as important as the benefits they provide.
- Vendor support needs to be considered as evaluating factor. This can be ascertained by talking to other customers and searching the Internet for testimonial information.
- Though not a requirement, transparency mode is useful when looking at a WAN Optimization standard. This allows the IP addresses of the actual server/client will be preserved instead of the traffic being tunneled using the IP addresses of the Wan optimizers.
- WAN monitoring and reporting as well as QOS schemes can run in the same manner before the installation of WAN Optimizers. Again, this is not required because most optimizers provide a method of doing these things but it should be considered.

IMPLEMENTATION GUIDANCE

Capacity will need to be determined by calculating the required concurrent number of workstation connections to the Data Center that needs acceleration. Without Exchange, recommended estimations have been identified 5 to 10 connections per workstation plus 15% growth as a guide to determine necessary capacity at the Data Center.

Many factors need to be considered while deploying and maintaining WAN optimizers. Many applications evolve or are upgraded on the client and server systems and this can have a detrimental effect on the method and effectiveness of the way the WAN optimizer operates. For instance, when MS Office was upgraded to 2007, encryption was turned on in all client systems in regards to their connections to the Exchange server. Most optimizers cannot effectively accelerate or compress encrypted or previously compressed data going across the wire. The staff that is responsible for the operations of the Optimizers must participate in all groups involved in changing profiles and applications on servers and clients. SMB signing and SMBv2 is another example of data types

that may cause issues. It will be important to use WAN optimizers that can handle these kinds of security requirements and changes.

4 TAXONOMY OF STANDARDS

ID	Primary Attribute	Secondary Attribute
1	Platform Chassis	Type (Modular vs. Static)
		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	Platform Packet Processor	Type (Upgradable)
		Packet processor Redundancy
		Layer 3 Routing Throughput (pps)
		Online Insertion/Removable Hot-Pluggable Support
		Memory (DRAM/Flash)
3	Platform Interface Cards	Interface Speed
		Interface Media Flexibility Support
		Queuing Properties
		Online Insertion/Removable Hot-Pluggable Support
4	Platform Operating System	Type (Modular)
		Features Supported
		Routing Protocols
		Quality of Service
		Protocols Supported
		Network Management
5	Platform Supportability	Technical Support

		Parts replacement
6	Platform Maturity	Operating System Maturity
		Hardware Maturity
		Quality Assurance Certification
7	WAN	Latency
		Transport Method
		Reliability/SLA/Packet Loss
		Path Diversity (Includes Local Loop)
		Carrier Diversity
		Capacity
		Quality of Service
		Network Carriers
		Traffic Analysis / Probes
9	Network Management	Out of Band Management
		Capacity Planning
		Network Management Protocols
		AAA
		Throughput
10	WAN Encryption	FIPS 140-2 Certification
		Throughput
		Hardware vs Software
		Latency
		Performed using Encapsulation(Tunneling) vs. Payload
		Supported Connection Count
		Cryptography
		Jumbo Frame Support
		Interface Type
		Bypass Capability
11	WAN Acceleration	Capacity
		Interface Type

		Implementation (inline vs. redirection)
		Redundancy
		Traffic Types Supported
		Centralized Management and Reporting
		Connection Management / Visibility
		Effectiveness (Application Acceleration, compression, Auto-Discovery)

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

Click here to enter additional acronyms.

APPENDIX D – CONTRIBUTORS

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

Name	Organization	Role
Adames, Juan	Region 3	
Ashe, David	Region 1	
Benson, Gary	Region 1	
Branch, Gary	Region 4	
Carpenter, Marc	EIE	
Cunagin, Mike	Region 1	
Dahl, Tim	EIE	
DeLong, Andrew	EIE	
Dugan, Mike	Region 2	
Finn, Mickey	EIE	
Haislip, Hal	Region 2	
Hina, John	FOD	
Julian, Michael	EIE	
Mehl, John	FOD	
Kinn, Eric	Region 4	
Moser, Robert	Region 1	
O'Brien, Casey	CDCO	
Paul, Tony	Region 5	
Trey, Marty	EIE	

APPENDIX E – MPLS SLA

Service	Key Performance Indicator (KPI)	Service Level	Performance Standards	Acceptable Quality Level	Penalty (Incident -Based Credits - are based on a single occurrence of service delivery)
Network-Based IP VPN	Availability (VPN)	Routine	99.90%	≥ 99.9%	No usage charge should apply during the period of the outage. The MRC for a service that was interrupted by an outage shall be prorated if the outage duration exceeded 12 minutes and was less than the duration required to qualify for a Time to Restore credit. This credit does not apply if the TTR credit is applicable. The amount of credit due for eligible service outages shall be calculated as follows: Credit =(MRC for the service that experienced the outage x (.025) x (duration of the outage in hours and tenths of an hour minus 0.2 hours (12 minutes)))
		Critical	99.99%	≥ 99.99%	
	Latency (CONUS) (Note 1)	Routine	70 ms	≤ 70 ms	Best Effort
	Latency (OCONUS) (Note 2)	Routine	150 ms	≤ 150 ms	Best Effort
	Time to Restore	With Dispatch (Note 4)	8 hours	≤ 8 hours	Customer is entitled to 50% of the MRC for the service, unless failure was due to documented delays caused by customer
		Without Dispatch (Note 3)	4 hours	≤ 4 hours	
	Provisioning	Routine	45 days		Credit is equal to 50% of the Non-Recurring Charge(s) (NRC(s)) or 50% of the MRC(s) for the entire order, whichever is greater, unless the failure to meet this performance objective was due to documented delays by the customer
		Critical	23 day		
	Disconnect	All	30 days		Best Effort

Notes:

1. Latency value is the average round trip transmission between Agency premise routers for an IP VPN with all of its CONUS sites. Latency metric does not apply for DSL, Cable High Speed, Wireless, and Satellite access methods. Relevant standards are RFC 1242 and RFC 2285. The contractor may propose to the Government more cost effective test and

measurement technique alternatives that meet or exceed the requirements in RFC 1242 and RFC 2285.

2. Latency value is the average round trip transmission between Agency premise routers for an IP VPN with its CONUS and OCONUS sites. Latency metric does not apply for DSL, Cable High Speed, Wireless, and Satellite access methods. Relevant standards are RFC 1242 and RFC 2285. The contractor may propose to the Government more cost effective

test and measurement technique alternatives that meet or exceed the requirements in RFC 1242 and RFC 2285.

3. VPN availability is measured end-to-end and calculated as a percentage of the total reporting interval time that the VPN is operationally available to the Agency. Availability is computed by the standard formula:

$$Av(VPN) = \frac{RI(HR) - COT(HR)}{RI(HR)} \times 100$$