United States

Department of Veterans Affairs

Design Document
for
OneVA Network Migration to a Dual Carrier MPLS Transport

Version 1.1

April 2007

Mickey Finn
Telecommunications Operations Management Service (005N2)
Office of Information and Technology

Table of Contents

1 Executive Summary 5
2 Design Document Introduction 6
2.1 Design Document Purpose 6
2.2 Document Scope 6
2.3 Document Definitions and Acronyms 6
3 DCMT Detailed Design 7
3.1 Current VA Network Design 7
3.2 DCMT Technical Requirements 7
3.2.1 Protocol Support 7
3.2.2 Capacity 7
3.2.3 IP Quality of Service 7
3.2.4 System Management 7
3.2.5 Security 7
3.3 DCMT End State Design 8
3.3.1 MPLS Design Constraints 9
3.3.2 IP Addressing and Autonomous System Numbers 10
3.3.3 IP Routing 11
3.3.4 DCMT Network and Systems Management 20
3.3.5 DCMT Security 21
4 DCMT Implementation 22
4.1 Phase I - Emulation of Current Core over Sprint PIP 22
4.2 Phase II - Parallel Operations of ATM and MPLS Carrier Clouds 22
4.2.1 PhaseII.1 Pilot Site Sprint PIP Deployment 22
4.2.2 PhaseII.2 Router Deployment 23
4.2.3 PhaseII.3 Out of Band Deployment 23
4.2.4 PhaseII.4 Circuit Turn Up and Testing 23
4.2.5 PhaseII.5 MPLS to MPLS Routing Enablement 23
4.2.6 PhaseII.6 ATM Routing Turndown 23
4.3 Phase III- MPLS Only Operation 23
5 DCMT Test Plan 24
5.1 Requirements Traceability Matrix 24
5.1.1 PhaseII.1 Pilot Site Sprint PIP Deployment 24
5.2 PhaseII.1 Pilot Site Sprint PIP Deployment 25
5.3 PhaseII.2 Router Deployment 26
5.4 PhaseII.3 Out of Band Deployment 26
5.5 PhaseII.4 Circuit Turn Up and Testing 27
5.6 PhaseII.5 MPLS to MPLS Routing Enablement 28
5.7 Phase II Migration Complete 29
6 System Operation 31
6.1 Initial DCMT Operational Handoff Plan 31
6.1.1 NMS Oversight 31
6.1.2 Capacity Planning 31

7 Adjunct Project Design 32
7.1 Regional / RDP MPLS Integration 32
7.2 WAN Encryption 32
7.3 IPv6 32
8 Appendix A – Rack Level Diagrams 35
9 Appendix B – Detailed router Configurations 37
9.1 Initial State Configuration 37
9.2 Final State Configuration 38
10 APPENDIX C – Detailed Hardware and Software BOM 48
10.1 Router BOM 48
10.2 Out of Band BOM 48

Revision History		
Date	**Author**	**Changes Made**
02/15/2007	Andrew DeLong	Initial Document

1 EXECUTIVE SUMMARY

The OneVA enterprise WAN network, as deployed by TMP, has been evolving over the past several years with the major goal of migrating all VA business units under a single enterprise network topology. At the completion of Phase IV, with the exception of NCA, all of the various VA business units have been merged into this methodology.

As successful as the TMP project was, there has also been a significant increase in network demands made to the backbone as well as the various regions that are straining this architecture. Major VHA enterprise projects, OCIS initiatives, RDPC deployments, and normal growth patterns are continuing to place serious demands on the scale and technical capabilities of the OneVA Enterprise network. Over the last 18 months, the backbone utilization has continued to grow at 6% per month, or 108% over the last 18 months. With new demands for telemedicine, remote support for radiology, convergence of voice and video, and the continuation of systems centralization to the RDPs, the capacity management team projects growth trends over the next 18 months to meet or exceed those seen to date.

Pursuant to the demands outlined above, the engineering team at OI&T headquarters has been working on several design recommendations that can migrate a portion of the present architecture to the next generation of WAN services. Based on earlier draft designs and analysis, the engineering team has decided to move to a dual carrier, MPLS WAN backbone architecture. By migrating to this new architecture across the current Distribution layer of the OneVA enterprise, this new backbone design will increase the scalability, functionality, and performance of the present WAN topology, while maintaining prudent controls over deployment and operational cost.

The project will be implemented in three phases to ensure that ongoing operations will be not be impacted. The first phase will use the new transport at the core sites, the second will deploy the new transport to all Distribution nodes, and the final stage will involve the decommissioning of the current ATM transport at those sites. As the new transport is enabled at the existing sites, those sites will begin to use the new circuits and hardware for a portion of their traffic while continuing to have the ability to use all facets of the current ATM WAN. The gradual implementation will minimize project risk while providing almost immediate capacity increase benefits.

2 DESIGN DOCUMENT INTRODUCTION

2.1 Design Document Purpose

This document is intended to capture the motivation for, implementation of, and desired results for the DCMT project.

2.2 Document Scope

The document is intended to cover the following areas:

- General Project Description
- Project Requirements
- Implementation methodology
- Security Mitigation plans
- Testing methodology

The following are outside document scope and documented elsewhere:

- Project Plan
- Project Implementation Logistics
- Security Plan

2.3 Document Definitions and Acronyms

Definitions

Acronyms

3 DCMT DETAILED DESIGN

3.1 Current VA Network Design

The current OneVA Network Transport system is described in the *TMP Phase III Design Document*. That system is based on an industry standard hierarchical WAN model utilizing a core layer, distribution layer, and access layer. The current system ties into the National Gateways for Internet and other External connectivity. A system Certification and Accreditation was performed on that system and will be updated and resubmitted with the addition of the new transport.

3.2 DCMT Technical Requirements

The following subsections detail the technical requirements of the new transport and systems used to support that transport. Each requirement will be tracked in the implementation phase to ensure it has been met or any shortfalls have been mitigated. The C&A and additional Inspector General Reports have pointed to general and specific risk areas in the VA Enterprise Network that need attention. Mitigation of those risks is a high priority of the DCMT project.

3.2.1 Protocol Support

REQ.1 Provide support for end to end support for IPv4
REQ.2 Provide mechanisms for IPv6 transport

3.2.2 Capacity

REQ.3 Provide at a minimum 50% more useable bandwidth for each Region or Data Center

3.2.3 IP Quality of Service

REQ.4 Provide a way to prioritize IP traffic into a minimum of four levels of service

3.2.4 System Management

3.2.4.1 Network Management Systems

REQ.5 Provide at a minimum the same level of Network instrumentation as currently implemented.

3.2.4.2 Out of Band Access

REQ.6 Provide a secure solution for all remote access
REQ.7 Provide a remote access solution that functions when normal in-band access is unavailable

3.2.5 Security

REQ.8 Provide at a minimum the same levels of access and logging that the current system employees
REQ.9 Provide an enhanced in-band access solution with enhanced security
REQ.10 Provide a VA controlled user authentication solution for the system

REQ.11 Provide a mechanism to allow for encryption of sensitive data being transmitted over the WAN

3.3 DCMT End State Design

The goal of the new VA routing architecture is to utilize all available bandwidth as efficiently as possible and to standardize router configurations as much as possible. Mutiprotocol Label Switching (MPLS) delivers highly scalable, differentiated, end-to-end IP services. The VA will utilize the MPLS VPN services from two providers: Sprint (Peerless IP (PIP) service offering) and AT&T (Private MPLS Network Transport offering).

The DCMT project will replace core and distribution routers through a lease package with AT&T. The new routers will replace the existing Sprint leased Cisco 7606 routers in the current core and the older Cisco 7200 routers at the distribution sites. Each distribution router will now have two WAN connections, one to each carrier network. Along with the reduced latency and Quality of Service (QoS) support MPLS should bring, the project will essentially double the available backbone bandwidth for the enterprise at each location. The existing dual distribution layer design will remain in place. The second pair of distribution routers will continue to provide the region full policy control for internal routing and policies.

Routers will exchange IP routes with each of the two service provider networks from each location. The VA traffic will be sent to the carriers and the AT&T and Sprint MPLS networks will provide the necessary transport mechanism to get the traffic to the appropriate destination. As a result of the new MPLS networks, we will no longer have a hierarchical network at the Core and Distribution. Traffic between Regions / Autonomous Systems (AS) will no longer need to traverse the Core. With the new MPLS implementation, all region-to-region traffic will flow directly from One-VA distribution router to One-VA distribution router. The four core locations will remain in Reston, Lee's Summit, Richardson and Santa Clara as Internet and Video Conferencing Gateways.

All Region routes will be advertised from each distribution node to each of the carrier MPLS networks with equal metrics. The VA Regions will split their networks such that one-half of their routes will prefer one of the distribution nodes while the remaining routes will prefer the second Regional One-VA distribution node. When the Region routes enter each of the MPLS networks (i.e. at the ingress of the provider network), the carrier networks will send destination traffic back to the Region through the appropriate distribution node. Traffic from the MPLS networks will also be split accordingly between the two Distribution nodes for each Region.

Redundancy is provided with multiple exit points out of a Region and dual MPLS network connections at each One-VA distribution node. Below is a diagram illustrating a typical Region connection to the MPLS networks along with the four Gateway locations.

Figure x. – MPLS Transport End State

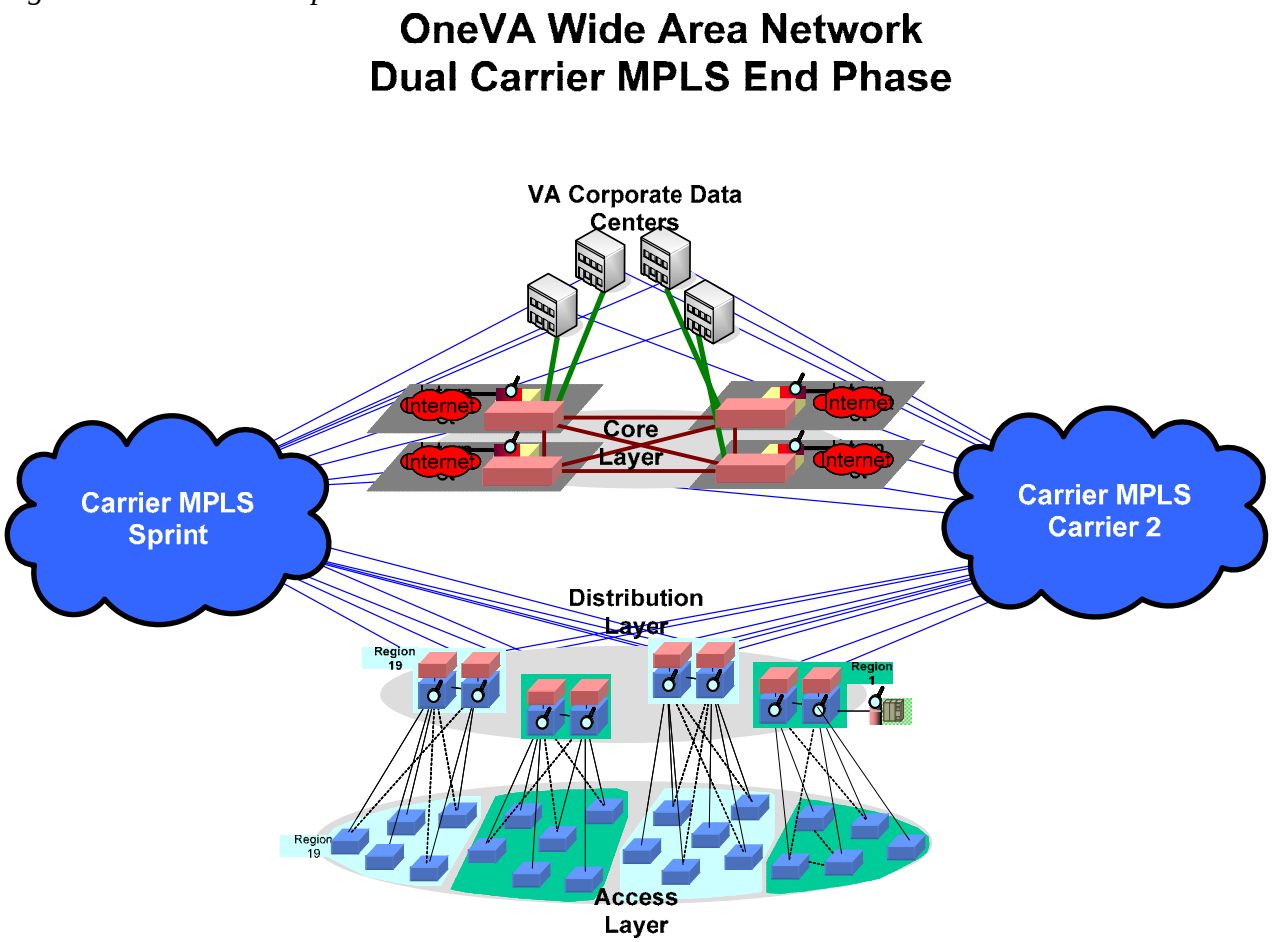

3.3.1 MPLS Design Constraints

As noted there are a number of benefits to the new MPLS service offerings from the carriers. However, there are also some difficulties introduced by the way the new technology is being implemented by those carriers. The primary constraint introduced is the limited support for, or non-existence, of equal cost routing within the carrier MPLS IP clouds. The carrier clouds will in general install a single best path for each destination. The current WAN relies on equal cost routing to provide load balancing as well as site fail-over. Additionally, routing to the Internet relies on the ability to have multiple default routes available on the WAN.

The Carriers' MPLS networks handle traffic and routing in different ways. AT&T provides MPLS services with load sharing features for up to two (2) equal routes. Sprint does not have any features enabled for equal cost paths, therefore, in order to standardize our router configurations we have designed the routing to accommodate unavailability of multi-path routing inside the PIP network. In order to utilize the two carrier clouds in the same way and keep the

overall architecture as simple and supportable as possible, the equal pathing available under AT&T will not be implemented at this time.

3.3.2 IP Addressing and Autonomous System Numbers

The VA will assign new blocks for the 10.x network to bring the DCMT Network online. The new transport will be addressed differently for each carrier:

3.3.2.1 Autonomous System Numbers

The VA Enterprise will continue to use the current AS numbering scheme for the existing regions as defined in the TMP Phase IV Document. The following additions will be made:

Site	**AS #**
RDP Sacramento	72
RDP Denver	73
RDP Brooklyn	74
RDP Philadelphia	75
Current Core Reston	92
Current Core Lee's Summit	93
Current Core Richardson	94
Current Core Santa Clara	95

3.3.2.2 Sprint PIP

The Sprint PIP Cloud will use the VA prescribed 10.net addressing and will follow the same scheme used for the current ATM cloud. The GRE tunnel interface will be assigned out of the same space using a similar scheme that will allow identification by region.

Specifically the IP schema will be:

Distribution Node1 CE<->PE - 10.250.region#.0/30
Distribution Node1 Tunnel to Gateway - 10.250.region#.100/30
Distribution Node2 CE<->PE - 10.250.region#.128/30
Distribution Node2 Tunnel to Gateway - 10.250.region#.200/30

Example - Region 5:
Martinsburg AS5: 10.250.5.2/30, Sprint PE Router 10.250.5.1/30
Martinsburg AS5 Tunnel: 10.250.5.102/30, Gateway Router 10.250.5.101/30
Silver Spring AS5: 10.250.5.130/30, Sprint PE Router 10.250.5.129/30
Silver Spring AS5 Tunnel: 10.250.5.202/30, Gateway Router 10.250.5.201/30

3.3.2.3 AT&T

The AT&T Cloud will be addressed using AT&T assigned IP Addresses. Those addresses are assumed to be in the 12.x.x.x IP range, but will be assigned by the carrier on a site by site basis

Note: The subnets used will be visible in the VA Routing Table in order to manage the interfaces involved.

3.3.2.4 Core/Gateway Router

The core routers will have a VLAN connection used for communications between the two routers which will be addressed in the 10.250.9x.0/29 space.

The core routers will have an additional interface, loopback1, used for termination of the GRE Internet and IPv6 tunnels. The IP will be carved out of the 10.250.9x.100/32 space allocated for the core router to core router communication.

3.3.2.5 Router and Distribution Node Device Management IP Space

Loopback IP Space will be assigned out of the current 10.254.x.x Management IP Space. The loopback 0 interface will be 10.254.100.reg#/32 and 10.254.100.1reg#/32 space.

A new network will be created for quasi out of band management of devices at the Distribution nodes. The network will be hosted directly on the MPLS router to ensure that it can be routed when there are network or IPS issues.

The IP Space used will be carved out of the 10.250.100.0 – 10.250.195 and be subnetted as /25 networks

- Router Management interface FA2/0 = .1 or .129
- IPS =.10 or .130
- NetScout =.20 or .140
- Local Switch Management =.30 or .150

Example - Region 5 devices:

Martinsburg Distribution Node Router FA2/0: 10.250.105.1/25
Martinsburg Distribution Node Router lo0: 10.254.100.5/32
Martinsburg IPS: 10.250.105.10/25
Martinsburg NetScout: 10.250.105.20/25
Martinsburg Mgt. Switch: 10.250.105.30/25
Silver Spring Distribution Node Router FA2/0: 10.250.105.129/25
Silver Spring Distribution Node Router lo0: 10.254.100.105/32
Silver Spring IPS: 10.250.105.130/25
Silver Spring NetScout: 10.250.105.140/25
Silver Spring Mgt. Switch: 10.250.105.150/25

3.3.3 IP Routing

3.3.3.1 Standard Routing Topology

A 2-port DS3 adaptor (PA-2T3+) will be used to terminate the DS3 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 both MPLS networks. 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 Sprint AS 1802). 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 to the networks that belong to that region 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. 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.

Every 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 1802 for Sprint 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. Based on lab testing conducted by the engineering team, the configuration command *bgp bestpath as-path multpath-relax* works by installing equal paths from differing AS into the Routing Information Base (RIB) which 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.

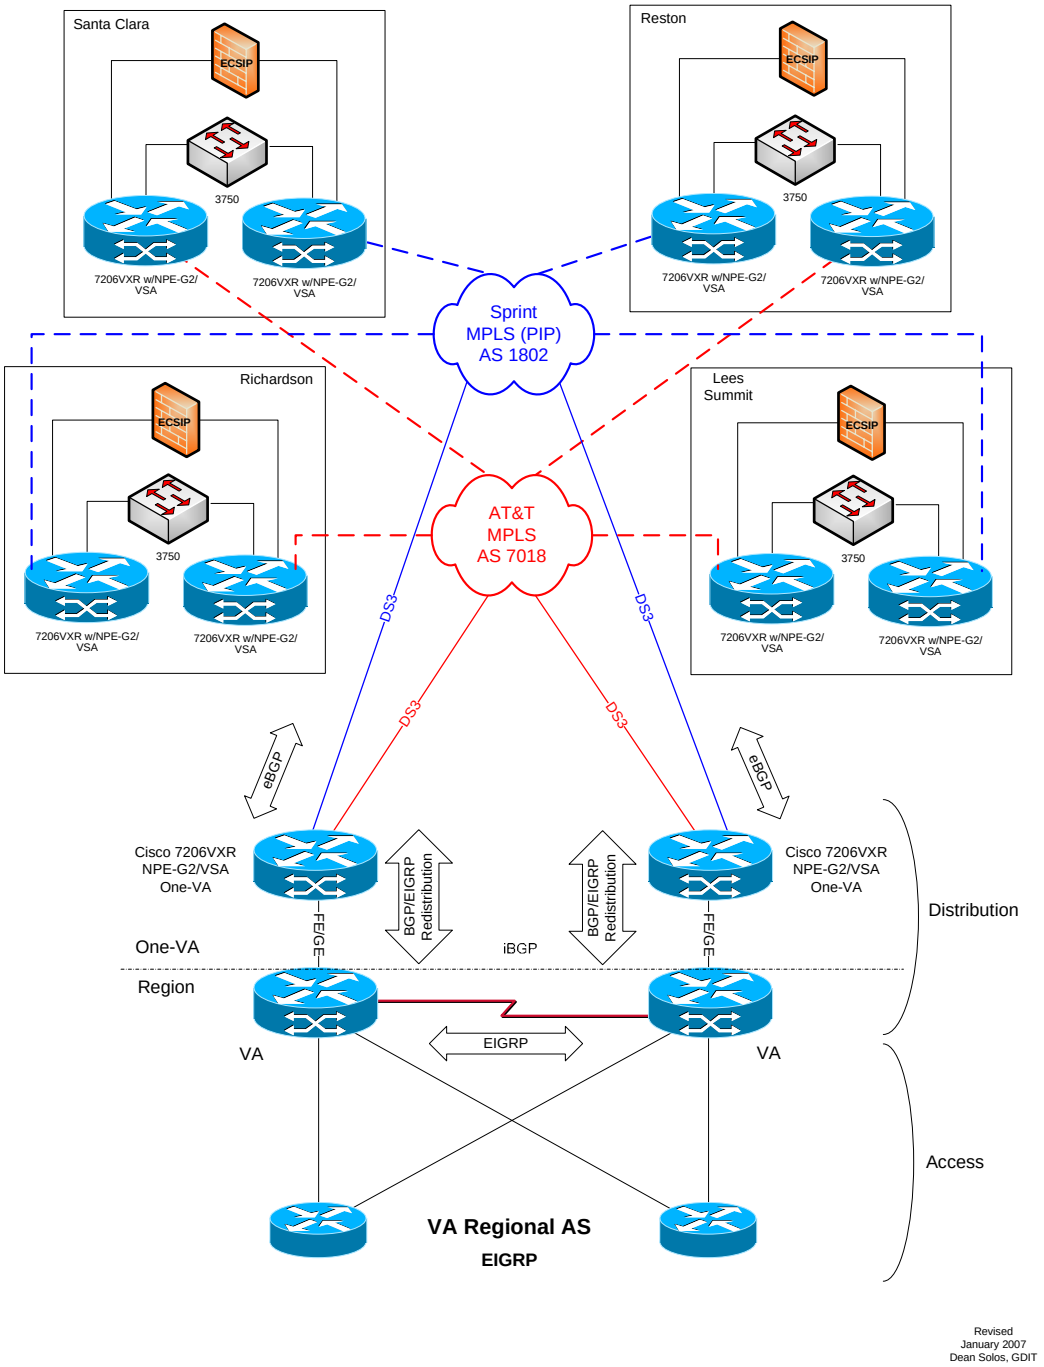

3.3.3.2 Default Routing and Internet Access

The existing dual distribution layer design will remain in place therefore, redundancy is provided with multiple exit points out of a Region and dual MPLS network connections at each One-VA distribution node. The second pair of VA Regional Distribution routers will continue to provide the region full policy control for internal routing and policies. We are redistributing the BGP routes into EIGRP.

The default route (0.0.0.0/0) will continue to be received from ECSIP as is currently the case. However, due to the limitations inherent in the MPLS networks with installing multiple equal cost paths, we will need to create tunnels to send the default route to an AS in a controlled manor:

- Tunnels will be configured using the Loopback IP addresses of the One-VA distribution router and one of the two routers located at the Gateway locations;
- BGP peering will take place between the One-VA routers and the Gateway location routers via the Tunnel;
- Only the default route (0.0.0.0/0) will be allowed over the Tunnels to the Regions;
- Routes from the Region/One-VA Distribution will be denied to the Tunnel interface;
- We will continue utilizing the Community List tagging that is currently used for the default routes from the Gateway locations to the One-VA Distribution routers.
- The Tunnels will also be needed as a transport for IPv6

The tunnels will allow the default route from the Gateway locations to be advertised to the Regions. All default (i.e., Internet-bound) traffic out of the Regions will use the Tunnels to get to the Internet. All return traffic will use the MPLS networks back to the destination Region.

3.3.3.2.1 Tunnel Overlay

In order to facilitate the default routing architecture and until both MPLS providers can support the IPv6 protocol, a GRE tunnel overlay will be used. The tunnel overlay will consist of one (1) GRE tunnel per distribution router to a selected gateway location. The gateway locations will stay consistent with the ATM networks current core to distribution connections. The exception will be datacenters which currently have four (4) UBR connections. Those Data Centers will only have tunnels to two locations to keep with the standard of one tunnel per distribution router. Each tunnel will utilize a unique assigned network as indicated in the section 3.3.2.3 Device Management IP Space.

Each tunnel will be built from the loopback0 address of the distribution to a unique loopback1 address on one of the gateway 7200 routers in each location. The loopback1 address will be independent of the loopback0 used for the gateway routers administration. By using a separate loopback address at the gateways, we can recover the tunnels in the case of a catastrophic failure of the primary tunnel endpoint 7200 hardware by rebuilding the tunnels to the remaining 7200 in service using the same loopback information.

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. At that time, tunnels between each gateway location will need to be established to pass only IPv6 routing information, no IPv4 information will pass. Once the MPLS carriers can support the IPv6 protocol, the tunnels will be reconfigured to again allow only the default route for IPv4. The default IPv6 route will also be added to the tunnel design if all four (4) gateway locations can support that configuration.

3.3.3.2.2 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 to the distribution node. Since the Distribution nodes are split between gateway locations, the other regional distribution node will continue to receive the default and outbound traffic will shift to that distribution node.

In the case of a distribution node's loss of both MPLS providers, by default the tunnel would 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.

3.3.3.3 QoS Design

Quality of Service (QoS) will play an integral part in the new MPLS transport. The current VA WAN has QoS implemented to help support Video Teleconferencing and other latency sensitive applications. QoS can help make better use of link bandwidth. If a link is saturated with traffic, QoS can mitigate the consequences and help control the impact. However, QoS cannot compensate for the overall lack of bandwidth.

QoS policies will be configured on the One-VA Distribution, Core and Data Center (DC) routers. The QoS will be used to support programs such as PolyTrauma, VANTS, future IP Telephony services, and other applications deemed mission critical or latency sensitive by the VA.

The VA NSOC will continue to mark packets at the ingress of the LAN interface on the One-VA distribution routers using the specific diffserv code points (DSCP) as described below. Queuing will be configured outbound on the serial interfaces of the Core and Distribution routers. Queuing will also be configured on the LAN interfaces of the One-VA distribution routers.

The specific QoS policies described are based on the available CoS Package offering from AT&T. The Sprint MPLS network has the ability to implement the current VA QoS policies. However, AT&T offers 24 profiles based on four (4) classes of service. Each of the four classes of service has a different bandwidth percent allocation based on the profile chosen. The VA has chosen separate profiles for the core and distribution routers, as shown in the table below.

Profile Number	Location	Bandwidth Allocation (%)			
		CoS 1 (DSCP EF)	CoS 2 (DSCP AF31/32)	CoS 3 (DSCP AF21/22)	CoS 4 (DSCP 0)
8	Core	50	0	0	50
14	Distribution	20	32	24	24

QoS policies will differ based on the expected traffic patterns. The core (or gateway locations) will provide service for voice, video and internet bound traffic. Therefore, the policies class of service will mirror those specifications.

In order to keep all router configurations standard, all QoS policies will be the same for both the AT&T and the Sprint networks.

The following classes will be used for the new MPLS transport on the distribution routers:

voice-video

Classification	The voice portion is included as a place holder for all future IP Telephony/Voice traffic. The video will be used to support PolyTrauma, VANTS and other video applications.
Marking	EF
Policy	Priority queuing using 20% of link bandwidth.

mission-critical-mgmt

Classification	The mission critical portion will support application such as VBA BDN and NCA Gravelocator. The 'mgmt' portion will be used to support all network management system traffic such as SNMP, syslog, telnet or SSH to a network device.
Marking	AF31
Policy	Guaranteed bandwidth of 40%.

data-high

Classification	Data high includes all internal VA traffic such as email.
Marking	AF21
Policy	Guaranteed minimum bandwidth of 30%.

class-default

Classification	Default class for all non-internal or Internet traffic.
Marking	DSCP 0 (BE)
Policy	Best Effort traffic.

3.3.3.3.1 QoS Service Type Mapping

3.3.3.4 Special Case Sites

While deploying the dual carrier MPLS transit, several sites were reviewed and given consideration as special cases. The sites include the internet and video gateways, Regional Data Processing Centers, and locations where one or more AS exists in the same physical location.

In the locations where two or more AS regions exist it was determined redundant to continue with this configuration. In these cases the two AS regions will be combined on the newly deployed routers with specific QoS and other configurations used to protect each original AS region's traffic from the other as represented in Figure 4 of Appendix A.

3.3.3.4.1 Core Locations

With the new MPLS VPN based network being deployed, the core will be reducing its role with inter AS traffic for the VA enterprise network. As part of the deployment, we are positioning two 7206 NPE-G2 routers in replacement of the single 7606 to handle the remaining internet and video traffic that will pass through these sites. Currently ECSIP uses an active/passive gigabit Ethernet configuration on a single VLAN of the 7600 to pass internet traffic to the enterprise network. With the new router deployment, ECSIP will have two active gigabit Ethernet interfaces and require equal cost static routes back to each device. This will allow us to per flow load share all of the internet traffic from the gateways to the clients over both MPLS networks.

The following drawing shows that we will initially connect the new MPLS routers g0/3 interfaces with the current 7600 router on g1/2 and g2/2 one port for each router. Once we have established that the connection is stable, the current OC3 Sprint connection will be moved to its designated router. Once the AT&T circuit is delivered, it will be connected to the second 7206. At or close to the same time, we will move the current ECSIP gigabit connections from the 7600 ports g1/1 and g2/1 to port g0/1 on each router and change the static routing to ensure we have equal paths to either router. This will allow traffic to continue to flow properly to the ATM connected regions but leave a solid pathway to migrate internet traffic to the MPLS configuration.

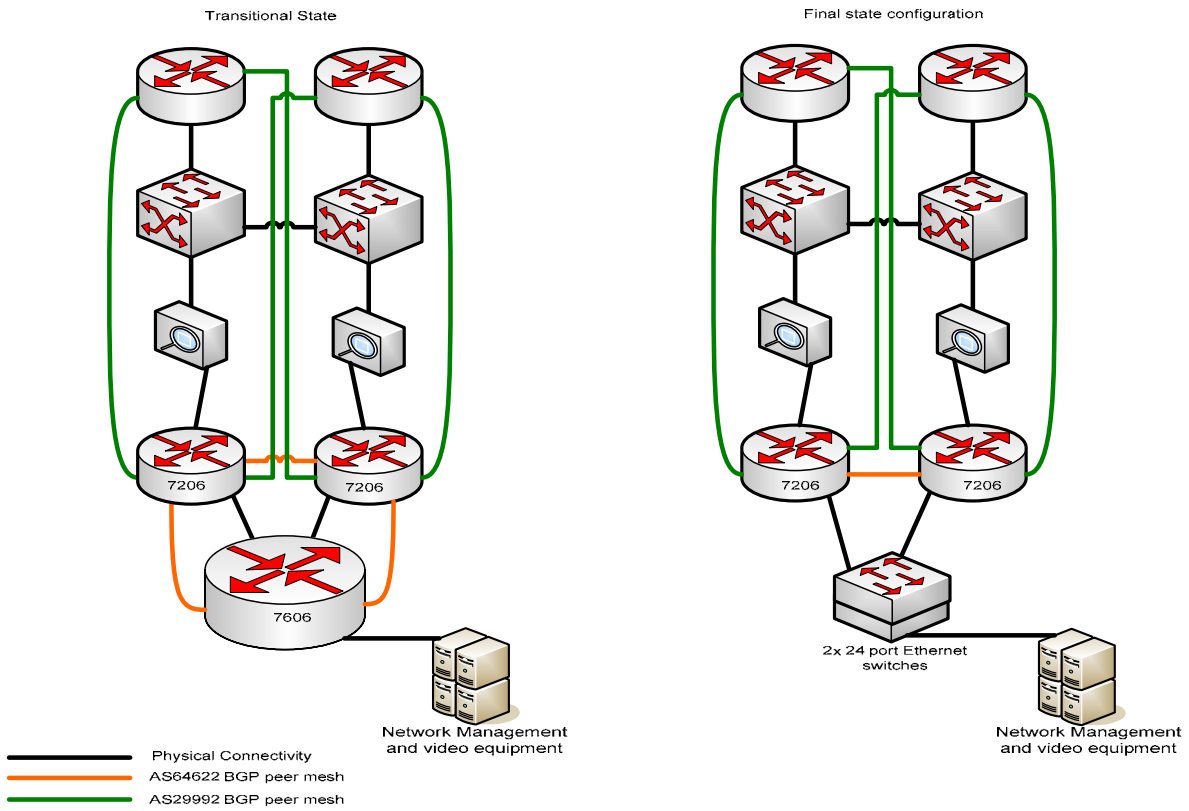

In preparation for the 7600 routers being removed, a pair of 24 port Ethernet switches will be deployed as the replacement for the 48 port 10/100 network card currently in use on the 7600. They will be connected to the 7206 routers on port g0/2 using GLBP and maintaining the current VLAN infrastructure. Global Load Balancing Protocol will allow use to run redundantly and utilize both links at the same time. GLBP will ensure we have the required levels of redundancy and allow us to continue to utilize both Ethernet ports in an active state.

Once the transition off the ATM network is completed, the 7600 routers will be removed from the network and BGP mesh. In the final state the traffic will flow out the ECSIP gateways and through either of the two 7206 routers and through the MPLS networks to their final destinations. To ensure after the transition that internet traffic is properly load sharing we will employ a restrictive tunnel configuration. The tunnels being used are to ensure internet requests go out the designated gateways and due to the gateways unique addressing will return from the same location. This solution works around the issues we had with the two networks handling of our routing tables and leaves us with the ability to use the tunnels for any IPv6 trials.

3.3.3.4.2 Silver Spring Data Center

The Silver Spring AS65025 (current Data Center) and AS65005 (Region 5) exist in the same physical location. The decision to combine the two networks will only be on the routers at the Silver Spring center side. The two regions will be collapsed into AS65005 and the router deployment will not mimic the standard regional deployment. The Martinsburg Distribution Node for AS5 will still receive a new 7206 and a DS3 to the two carriers. Both networks will

connect on individual Ethernet ports on the two 7206 routers and maintain the current EIGRP processes identifiers they use today.

The subnets will be announced differently to the carriers with region 5 being preferred over one of the DS3's and the data center preferred on the other DS3. Traffic out of both the region and data center will load balance over both of the DS3's. In order to protect the region and datacenter from overloading each other's traffic, a specific QoS configuration will be used on the DS3 interfaces. The configuration will guarantee 65% of the total bandwidth to the datacenter and the remaining 35% to the region. The route preferences will ensure that neither network can do overload the either for traffic heading to the Silver Spring location. Both networks will have failover to the OC3 without route preference.

3.3.3.4.3 Falling Waters Data Center

The Falling Waters data center AS65026 (current Data Center), AS65061 (VACO), and AS65124 (NSOC) exist in the same location. AS65061 and AS65026 will be combined into AS65026 due to concerns the NOC network will not be combined with the data center and VACO networks and will continue to exist as a separate AS. The datacenter will receive two 7206 routers and a DS3 to each MPLS carrier while VACO will receive a 7206 and a DS3 from each MPLS carrier for the Washington, DC distribution node.

The subnets will be announced differently to the carriers with VACO being preferred over one of the DS3's and the data center preferred on the other DS3. Traffic out of both the region and datacenter will load balance over both of the DS3's. In order to protect the region and datacenter from overloading each others traffic a specific QoS configuration will be used on the DS3 interfaces. The configuration will guarantee 65% of the total bandwidth to the datacenter and the remaining 35% to VACO. The route preferences will ensure that neither network can do overload the either for traffic heading to the Silver Spring location. Both networks will have failover to the DS3 without route preference.

3.3.3.4.4 Martinsburg VAMC, OIFO and NOC Lab

The Martinsburg VAMC and OIFO have already converged the LAN networks and no special changes are required. The old NOC location and the current temporary location of the Internet Gateway LAB will be integrated into the VAMC MPLS router. Exact engineering on the Ethernet link to the router location will be engineered with the help of the local LAN and Facility managers. If a separate physical path cannot be used, a separate VLAN will be requested and connected to its own gigabit port on the 7200 routing. The level of QoS protection between the two networks will need to be established due to the low utilization of the old NOC ATM link.

3.3.3.4.5 NCA Quantico and Culpepper

The NCA Quantico and Culpepper locations, while physically separated, are considered part of the same Data Center. Due to their status and internal network configuration, Quantico and Culpepper will be configured similar to a region. The only variation is that Culpepper is the backup location and only active during a failover. Routing information passed to the MPLS network will reflect this primary and secondary configuration, however all BGP communities sent from both locations will be identical so that routing information is consistent even if the NCA is in a failover configuration.

3.3.3.4.6 RDP Regions 1 and 4

The RDP Regions 1 and 4 data centers will also receive MPLS carrier connections. Both designs are based on two (2) data center locations within the regional areas. These locations, while somewhat symbiotic, will receive two routers in each location as if they are independent with on carrier connection per router equaling two connections per datacenter and four connections per region. The final network design including AS allocation, network route preferences and other engineering issues is still being discussed. The remaining engineering information will be completed at a later date.

3.3.4 DCMT Network and Systems Management

3.3.4.1 Out of Band Management System

The VA is deploying a more robust and broader scoped version of the Out of Band System in use at the VA today. That system is based on hardware and software products from CDI. The software system is comprised of a Server component (DDM) and a client component (SSM). The hardware system is comprised of secure dial out modems (Uniguard) and modem encryptors (CDI SSE) and end host terminal servers (PA-44). A complete BOM is attached as appendix x

3.3.4.1.1 Secure Session Manager - SSM

The Secure Session Manager (SSM) program works in conjunction with the database contained within CDI's Distributed Database Manager (DDM) program. There are two types of SSM programs. One that operates with Microsoft Access Database (ODBC) and one that operates with Structured Query Language Server 2000 (SQL). The SSM program allows a manager to access a local or remote CDI device in an encrypted mode and/or Authentication mode from the manager's PC. The encrypted keys are contained in the SSE device attached to the SSM PC's USB port. These encrypted keys have been previously assigned to the CDI devices by CDI's Distributed Database Manager DDM program. The SSM software will allow encrypted access to the CDI device but the manager must log on to that device before gaining access to the equipment connected to that device and must have a CDI SSE device attached to the SSM PC's USB port.

3.3.4.1.2 Distributed Database Manager - DDM

DDM Client for SQL Server 2000 Server version and the DDM for the ODBC version are both software products designed to remotely manage the databases of CDI's UniGuard, Port Authority, Port Authority SAM, MultiGuard and SSE devices. The number of devices that can be managed is virtually infinite, limited only by the storage capacity of the computer system on which DDM databases are installed. The program is designed to function using Windows 2000 or Windows XP operating systems. The VA is deploying the DDM SQL version which can handle multiple concurrent sessions. Through the SQL Server 2000 database, a number of DDM managers can access the Database simultaneously. The DDM databases are stored on the SQL Server 2000 database and are easily modified, synchronized and updated. All information can be backed up for easy and safe restoration in the event of system failure using the SQL Server 2000 database utilities. All communications with the system is encrypted for maximum security.

3.3.4.1.3 CDI Modems

The VA will deploy three types of modems based on function they are intended to perform:

- Uniguard Encrypted Modems – Used as a fixed dial out modem in the NSOC.
- CDI SSE – Used to encrypt a standard modem for dial out. These will be used by mobile support personnel and used in disaster recovery or COOP situations.
- PA-44 - Deployed to each physical WAN location in the Data Centers and Distribution Nodes. The system provides for remote management via four console ports as well as remote power cycle capability.

3.3.5 DCMT Security

The DCMT will leverage the security posture put in place by the TMP project. That project was Certified and Accredited in June, 2005. A complete review of the system security plan and systems controls will be performed.

3.3.5.1 Current One-VA Security Risks Mitigated by DCMT

More secure router management traffic – The current system relies on Telnet for a number of management connections to the infrastructure routers. The new routers are capable of using SSH for communications to NSOC workstations as well as Network Management Applications providing a more secure method of managing the infrastructure.

WAN encryption – The new platform was specified with a hardware accelerated IPSec encryption card and the platform is being FIPS 140-2 certified. The initial rollout will not incorporate WAN encryption, but the feature is ready to be deployed shortly after the completion of the DCMT. The current system does not have software features supporting encryption and the hardware platform has limited encryption capabilities.

VA control of infrastructure authentication – The new transport and router platform will be under VA control. This allows the VA to control and audit all authentication to the platform. The current authentication and authorization system is part of a leased service.

3.3.5.2 Security Risks Introduced

The following security risks are introduced by the new transport and will be documented and re-certified.

Second Carrier – The new transport introduces a new carrier, AT&T, and that carrier's systems. Those security controls identified in the original C&A specific to Sprint will need to be extended and tested with AT&T.

Layer 3 Exposure – The move to VPN based MPLS exposes an additional layer of the OSI Stack to the carriers transporting the VA data. The current network is ATM based and only exposes layer 2 traffic to the carrier. Although the exposure is roughly the same, layer 3 traffic is generally easier to hijack.

4 DCMT IMPLEMENTATION

4.1 Phase I - Emulation of Current Core over Sprint PIP

The first phase of the implementation will use tunnels to mimic the existing point-to-point ATM network running between the core locations. The tunnel overlay was chosen because it represented minimal change from the existing network topologies. The phase I network is comprised of four (4) OC3 Sprint PIP interfaces in the existing Core 7606 routers. Those interfaces exchange routing updates with the Sprint PIP could, but only the connected networks are injected into that Sprint PIP could. Six (6) IP GRE tunnels are built between the four core locations. The TMP routing architecture is then run over those six GRE tunnels instead of the legacy ATM PVPs between the core sites.

4.2 Phase II - Parallel Operations of ATM and MPLS Carrier Clouds

It is planned that for the majority of the project, both the legacy TMP transport and the new MPLS transport will operate concurrently. This allows for a gradual migration and an augmentation of bandwidth. However, it does increase the overall complexity of the VA network during transition. Parallel operations will be deployed in phases to mitigate risk to the VA Enterprise.

4.2.1 PhaseII.1 Pilot Site Sprint PIP Deployment

The first deployment will occur to the following beta test sites:

- Philadelphia DC (AS11)
- Hines DC (AS13)
- Little Rock (AS16)
- NSOC (AS124)

The sites will use the legacy ATM router with a single port PA-T3 serial adapter added to terminate the new Sprint PIP circuits. Those sites will be brought up on PIP, exchanging the same route table as they currently send to the ATM core network.

As sites join the new Sprint PIP cloud they will propagate the PIP route table with the routes they own. There is a specific route map that will prohibit them from sending routes not associated with the AS they belong too.

As they join the new Sprint PIP cloud, the sites will start to use that cloud for a limited amount of traffic. Specifically, they will use the Sprint PIP cloud to talk to the other PhaseII.1 Pilot sites. This segregation is done by natural BGP best path selection and by assignment of metrics. AS's that join the Sprint PIP cloud will see legacy ATM connected AS's as a shorter path in all cases except for those other sites on that Sprint PIP cloud. If two sites are on the Sprint PIP cloud they will see each other with equal BGP AS paths on both the ATM and Sprint PIP cloud. A route map will set metrics based on a BGP community parameter that will cause the Sprint PIP cloud routes to be preferred.

4.2.2 PhaseII.2 Router Deployment

The new MPLS routers will be deployed to all Distribution, Core and Data Center sites. The initial configuration will be a generic configuration that will allow for remote connectivity and testing of the new MPLS circuits being delivered. The router BOM can be found in *Appendix C*. The base configuration and final configurations can be found in *Appendix B.*

The routers will be shipped with the base configuration from the NSOC to each location. The site POCs will rack and power the new routers.

4.2.3 PhaseII.3 Out of Band Deployment

The OOB system will be deployed at each location as well as the Falling Waters and Hines NSOCs. The system will be configured and installed by a Sprint CPE technician. That technician will work with the NOC to replace the existing production OOB modem onsite. Additionally, the OOB solution will be connected to the production IPS devices to provide the NSOC with remote capabilities for those devices. Both interim and final state configurations for the OOB system are detailed in *Appendix A.*

4.2.4 PhaseII.4 Circuit Turn Up and Testing

The new MPLS circuits will be provisioned and extended by the carriers and will land on the new MPLS router. The circuits will be turned up and tested initially by addressing them with those IP Addresses defined in section 3.3.2 and by PING'ing between the MPLS router (CE) and the carrier MPLS router (PE).

4.2.5 PhaseII.5 MPLS to MPLS Routing Enablement

PhaseII.5 will closely follow the PhaseII.1 methodology. The sites will be coordinated with and will be cut over one AS at a time. The traffic for that AS will exhibit the same behavior as described in PhaseII.1 and will start communicating with other AS's that have been cut to the new MPLS cloud.

4.2.6 PhaseII.6 ATM Routing Turndown

When a large majority of the existing sites have been fully cut to the PhaseII.5 MPLS architecture the ATM circuits serving those sites may be turned down.

The exact percentage of sites that must be cut over before ATM turndown is not known but assumed to be approximately 70%. Caution must be exercised in moving off the ATM transport early. Until all sites are transitioned, the existing core sites will have to act as a gateway between the legacy ATM and new MPLS clouds. Pushing too much traffic through that gateway point could artificially saturate the OC3 links at the existing core locations.

4.3 Phase III - MPLS Only Operation

The end state goal of the MPLS transport migration is to decommission the existing ATM network at the Core and Distribution.

5 DCMT TEST PLAN

The DCMT system will be tested as each phase is deployed as well as once phased deployment is complete.

5.1 Requirements Traceability Matrix

The following tables capture testing that will be done to ensure the requirements as defined in *Section 3.1* are met.

REQ.1 Provide support for end to end support for IPv4
REQ.2 Provide mechanisms for IPv6 transport
REQ.3 Provide at a minimum 50% more useable bandwidth for each Region or Data Center
REQ.4 Provide a way to prioritize IP traffic into a minimum of four levels of service
REQ.5 Provide at a minimum the same level of Network instrumentation as currently implemented
REQ.6 Provide a secure solution for all remote access
REQ.7 Provide a remote access solution that functions when normal in-band access is unavailable
REQ.8 Provide at a minimum the same levels of access and logging that the current system employees
REQ.9 Provide an enhanced in-band access solution
REQ.10 Provide a VA controlled user authentication solution for the system
REQ.11 Provide a mechanism to allow for encryption of sensitive data being transmitted over the WAN

5.1.1 PhaseII.1 Pilot Site Sprint PIP Deployment

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING connectivity testing between each end of the point to point circuits	PING reply	PING and traffic
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	GRE tunnel administratively up and protocol up	Tunnel tested and in service
REQ.3	CAP	Confirm the reported usable bandwidth between sites	Minimum of 50% average increase across all circuits	Core sites moved from 60Mbps to 155Mbps, a 258% increase.
REQ.4	QoSQ	Show QoS queue structure exists	Show command demonstrates that the QoS queues are in place	Queues are in place and show traffic in each class
REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	Mapped devices and interfaces, confirmed to change status on maps	Interfaces mapped, test confirms state change.
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A
REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being	Syslog and SNMP Traps	Confirmed

		captured on NMS	being forwarded and collected at the NSOC	
REQ.9	SEC1	N/A, No change in this phase	N/A	N/A
REQ.10	AAA	N/A, No change in this phase	N/A	N/A
REQ.11	ENC	N/A, No change in this phase	N/A	N/A

5.2 PhaseII.1 Pilot Site Sprint PIP Deployment

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING connectivity testing between each end of the point to point circuits	PING reply	PING and traffic
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	N/A	N/A
REQ.3	CAP	Confirm the reported usable bandwidth between sites	Minimum of 50% average increase across all circuits	PIP added 45Mbps to each Node a 107% increase.
REQ.4	QoSQ	Show QoS queue structure exists	Show command demonstrates that the QoS queues are in place	Queues are in place and show traffic in each class
REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	Mapped devices and interfaces, confirmed to change status on maps	Interfaces mapped, test confirms state change.
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A
REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being captured on NMS	Syslog and SNMP Traps being forwarded and collected at the NSOC	Confirmed
REQ.9	SEC1	N/A, No change in this phase	N/A	N/A
REQ.10	AAA	N/A, No change in this	N/A	N/A

		phase		
REQ.11	ENC	N/A, No change in this phase	N/A	N/A

5.3 PhaseII.2 Router Deployment

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING connectivity testing between each end of the point to point circuits	N/A	N/A
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	N/A	N/A
REQ.3	CAP	Confirm the reported usable bandwidth between sites	N/A	N/A
REQ.4	QoSQ	Show QoS queue structure exists	N/A	N/A
REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	Mapped devices and interfaces, confirmed to change status on maps	Interfaces mapped, test confirms state change.
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A
REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being captured on NMS	Syslog and SNMP Traps being forwarded and collected at the NSOC	Confirmed
REQ.9	SEC1	N/A, No change in this phase	N/A	N/A
REQ.10	AAA	N/A, No change in this phase	N/A	N/A
REQ.11	ENC	Confirm new hardware is capable of providing encryption for the WAN		

5.4 PhaseII.3 Out of Band Deployment

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING	N/A	N/A

		connectivity testing between each end of the point to point circuits		
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	N/A	N/A
REQ.3	CAP	Confirm the reported usable bandwidth between sites	N/A	N/A
REQ.4	QoSQ	Show QoS queue structure exists	N/A	N/A
REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	N/A	N/A
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A
REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being captured on NMS	N/A	N/A
REQ.9	SEC1	N/A, No change in this phase	N/A	N/A
REQ.10	AAA	N/A, No change in this phase	N/A	N/A
REQ.11	ENC	N/A, No change in this phase	N/A	N/A

5.5 PhaseII.4 Circuit Turn Up and Testing

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING connectivity testing between each end of the point to point circuits	PING reply	PING and traffic
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	N/A	N/A
REQ.3	CAP	Confirm the reported usable bandwidth between sites	N/A	N/A
REQ.4	QoSQ	Show QoS queue structure exists	Show command demonstrates that the QoS queues are in place	Queues are in place and show traffic in each class

REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	Mapped devices and interfaces, confirmed to change status on maps	Interfaces mapped, test confirms state change.
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A
REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being captured on NMS	Syslog and SNMP Traps being forwarded and collected at the NSOC	Confirmed
REQ.9	SEC1			
REQ.10	AAA			
REQ.11	ENC	N/A, No change in this phase	N/A	N/A

5.6 PhaseII.5 MPLS to MPLS Routing Enablement

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING connectivity testing between each end of the point to point circuits	PING reply	PING and traffic
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	GRE tunnel administratively up and protocol up	Tunnel tested and in service
REQ.3	CAP	Confirm the reported usable bandwidth between sites	Minimum of 50% average increase across all circuits	
REQ.4	QoSQ	Show QoS queue structure exists	Show command demonstrates that the QoS queues are in place	Queues are in place and show traffic in each class
REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	Mapped devices and interfaces, confirmed to change status on maps	Interfaces mapped, test confirms state change.
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A

REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being captured on NMS	Syslog and SNMP Traps being forwarded and collected at the NSOC	Confirmed
REQ.9	SEC1	N/A, No change in this phase	N/A	N/A
REQ.10	AAA	N/A, No change in this phase	N/A	N/A
REQ.11	ENC	N/A, No change in this phase	N/A	N/A

5.7 Phase II Migration Complete

Requirement	Test Name	Test Description	Expected Result	Actual Result
REQ.1	PING	Standard IPv4 PING connectivity testing between each end of the point to point circuits	PING reply	PING and traffic
REQ.2	GRE	Ensure Generic Route Encapsulation is in place and ready for IPv6	GRE tunnel administratively up and protocol up	Tunnel tested and in service
REQ.3	CAP	Confirm the reported usable bandwidth between sites	Minimum of 50% average increase across all circuits	Core sites moved from 60Mbps to 155Mbps, a 258% increase.
REQ.4	QoSQ	Show QoS queue structure exists	Show command demonstrates that the QoS queues are in place	Queues are in place and show traffic in each class
REQ.5	NMS	Confirm all newly introduced elements are managed. Take element down and ensure it is reflected in NMS system	Mapped devices and interfaces, confirmed to change status on maps	Interfaces mapped, test confirms state change.
REQ.6	OOB1	N/A, No change in this phase	N/A	N/A
REQ.7	OOB2	N/A, No change in this phase	N/A	N/A
REQ.8	LOG	Confirm new element information is being captured on NMS	Syslog and SNMP Traps being forwarded	Confirmed

			and collected at the NSOC	
REQ.9	SEC1	N/A, No change in this phase	N/A	N/A
REQ.10	AAA	N/A, No change in this phase	N/A	N/A
REQ.11	ENC	N/A, No change in this phase	N/A	N/A

6 SYSTEM OPERATION

6.1 Initial DCMT Operational Handoff Plan

The Engineering team and the NSOC will work to update the Standard Operating Procedures (SOP). Those procedures will capture the interaction between the NSOC, the new transport, and the carriers providing that transport.

6.1.1 NMS Oversight

The new infrastructure will be incorporated in the existing NMS systems.

6.1.2 Capacity Planning

The new transport introduces a complexity to the capacity planning aspect of the network. The current network is instrumented with RMON probes on the physical interfaces. The new transport will be instrumented on the inbound Ethernet interface. Traffic on the physical interfaces will have to be collected via SNMP MIBs. Additionally, the extension of QoS to the WAN creates multiple logical pipes that will have to be managed as the current ATM PVPs are managed. The existing Capacity Planning SOP will be updated to capture the methods used to oversee the transport.

7 ADJUNCT PROJECT DESIGN

The following projects are greatly impacted by the DCMT design and deployment. The dependencies and

7.1 Regional / RDP MPLS Integration

The interaction with the RDPs will be further defined in the future.

7.2 WAN Encryption

The hardware being deployed with the DCMT is capable of encrypting approximately 300Mbps of traffic onto the WAN. As the requirements for encryption are defined the exact method of deploying that encryption will be further defined.

7.3 IPv6

As documented in the general routing and default routing sections of this document, a Generic Route Encryption (GRE) tunnel overlay will be introduced on the VA Enterprise backbone. The overlay is necessary for controlled propagation of default routing and as a mechanism for the IPv6 testing and transition. The carrier MPLS VPN networks we are deploying are not IPv6 aware at this time. They do not support peering using an IPv6 address nor do they support the propagation of IPv6 routes within the MPLS cloud. Exact timeframes and methods used for IPv6 support on the carrier MPLS clouds are not known at this juncture.

In order to accommodate the IPv6 transition and to allow for maximum flexibility in transition methodology the GRE tunnels will be introduced. These tunnels allow the VA to directly address each tunnel interface with a Native IPv6 Address. The GRE tunnels will support the use of an IPv4/IPv6 dual stack transition across the network backbone. The MPLS network being deployed does support other methods for IPv6 transition such as tunneling. Having both IPv6 transition methods available will allow the VA maximum flexibility in transition.

The VA will initially deploy an IPv6 overlay that will be capable of serving each VA site on the backbone with the IPv6 protocol. The overlay will be deployed at one of the East Coast and one of the West Coast Gateway locations. The overlay via the GRE tunnels will be capable of being extended to each region through one demarc location. The IPv6 infrastructure will then be extended to each VA site, such as a Medical Center, over the Regional infrastructure. Since that infrastructure varies, mechanisms used will vary by region. Most regions employ circuitry that supports native IPv6 transport. Hardware has been evaluated throughout the VA and those items that are not IPv6 capable are being replaced. The routing hardware being introduced by the DCMT project is fully IPv6 capable. The software and routing operating system is also IPv6 capable. The complete bill of materials for the routing platform is provided in Appendix x.

The Internet gateway locations at the Gateway location selected will be capable of transferring IPv6 packets to and from the Internet. The overall layout of those gateways is being further defined as is the exact path from the internal network to the IPv6 Internet. The VA is deploying connections at one or more Gateway locations to Internet2. That connection will allow the VA to announce their IPv6 block as well as collaborate with others via IPv6. The figure below captures the GRE tunnel and IPv6 Internet Gateway connectivity.

Figure 7.3-1 IPv6 GRE and ISP Connectivity

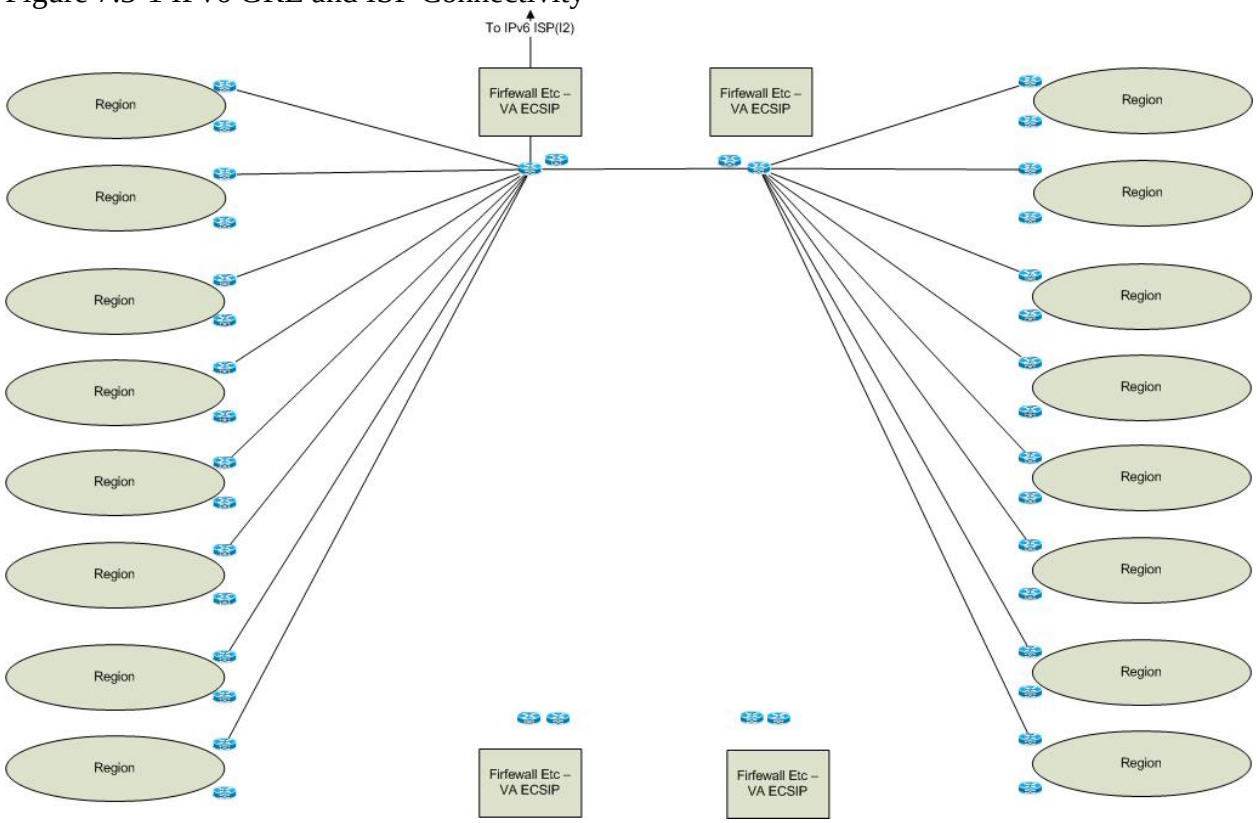

8 APPENDIX A – RACK LEVEL DIAGRAMS

Figure 1 – Distribution Node Diagram – PhaseII.3

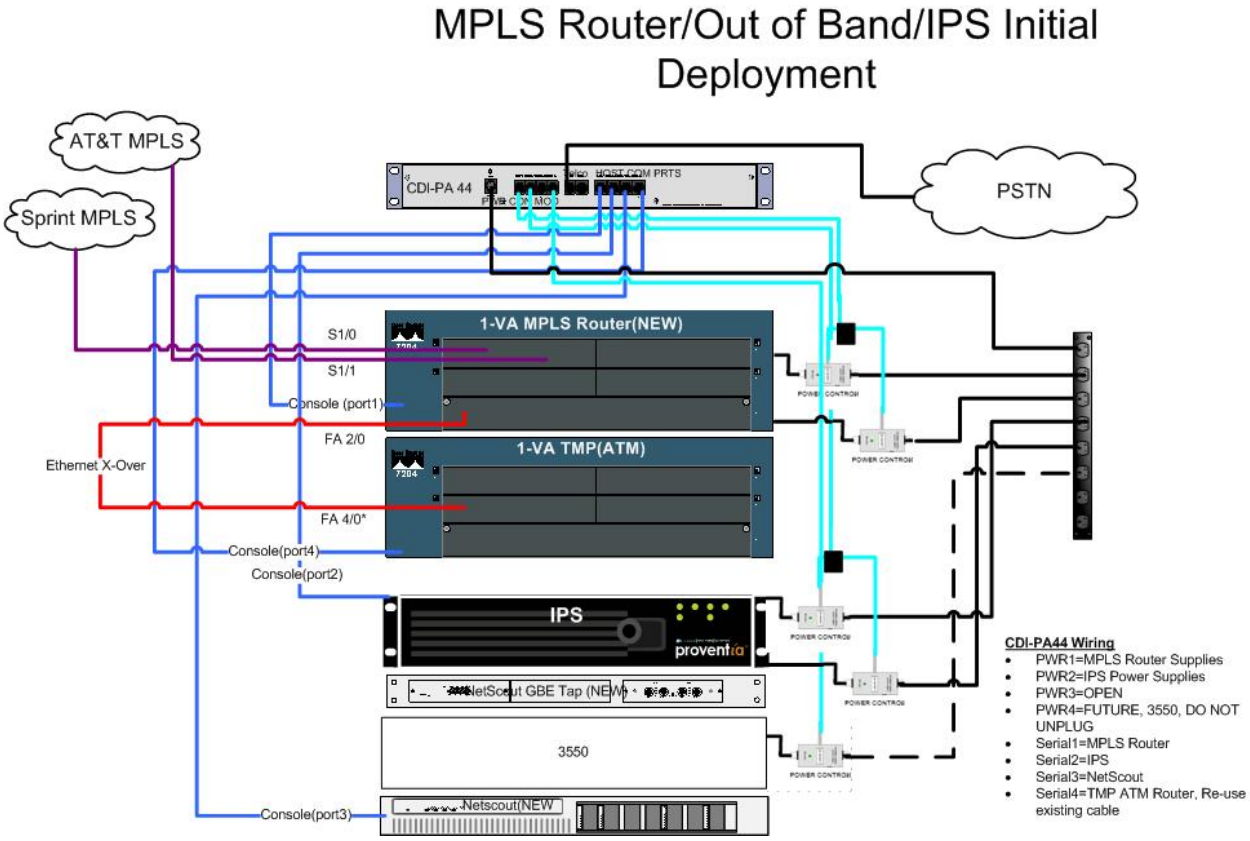

Figure 2 - Distribution Node Diagram – PhaseII.5

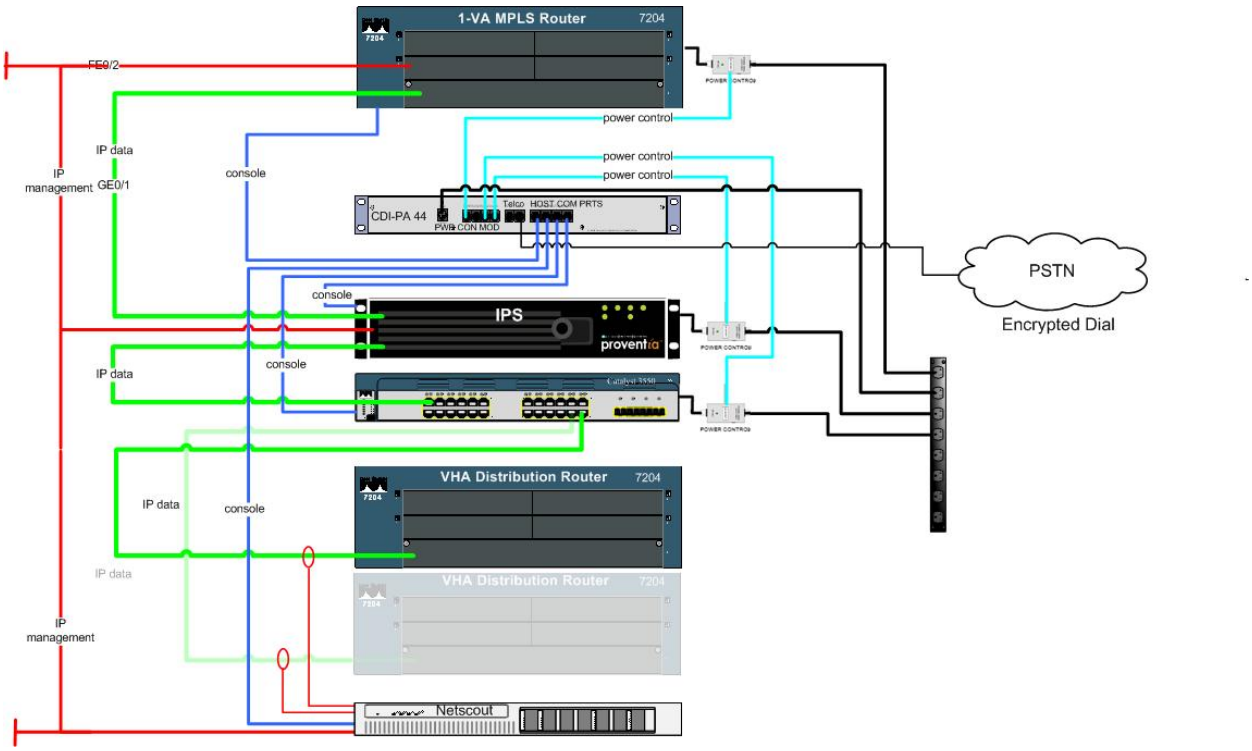

Figure 3 – Gateway Site Diagram

Figure 4 – Non Standard Distribution Sites

9 APPENDIX B – DETAILED ROUTER CONFIGURATIONS

9.1 Initial State Configuration

Revision 2/17/2007

```
!
!MPLS Baseline configuration
!
no service udp-small
no service tcp-small
service password
service timestamps debug datetime localtime show-timezone
service timestamps log datetime localtime show-timezone
!
hostname HOSTNAME
!
logging buffered 16000 debug
no logging console
enable secret t3sting
!
clock timezone EST -5
clock summer-time EDT recurring 2 Sun Mar 2:00 1 Sun Nov 2:00
ip subnet-zero
ip cef
ip classless
!
no ip domain lookup
!
ip domain-name net.va.gov
!
ip sla monitor responder
!
!
interface FastEthernet0/2
 ip address 192.168.1.1 255.255.255.252
 duplex full
 speed 100
 no shut
!
banner motd ^C

                        Warning

This system is for the use of authorized users only.  This system is
monitored to protect against its unauthorized use and for maintenance
purposes.  In the course of such monitoring the activities of
```

authorized users may be monitored and/or recorded by systems personnel.

Anyone accessing and using this system expressly consents to
system monitoring. If monitoring reveals evidence of unauthorized
or criminal activity, systems personnel may provide that evidence to law
enforcement officials.
^C
!
line vty 0 8
password 4Va4tmp4
!
!
end

9.2 Final State Configuration

Revision 2/17/2007

```
!
!MPLS Baseline configuration
!
no service udp-small
no service tcp-small
service password
service timestamps debug datetime localtime show-timezone
service timestamps log datetime localtime show-timezone
!
hostname AS61-Baseline-Test
!
logging buffered 16000 debug
no logging console
enable secret t3sting
!
aaa new-model
aaa authentication login default group tacacs+ enable
aaa accounting exec default start-stop group tacacs+
aaa accounting commands 15 default stop-only group tacacs+
!
aaa session-id common
clock timezone EST -5
clock summer-time EDT recurring 2 Sun Mar 2:00 1 Sun Nov 2:00
ip subnet-zero
ip cef
ip classless
!
no ip domain lookup
```

```
!
ip domain-name net.va.gov
!
ip sla monitor responder
!
class-map match-all in-data-high
description Internal Network Traffic
 match access-group name qos-data-high
!
class-map match-all in-mission-critical-mgmt
description Mission-Critical Data & Management
 match access-group name qos-mission-critical-mgmt
!
class-map match-all in-voice-video
description Voice/VoIP/IPT/video
 match access-group name qos-voice-video
!
class-map match-all out-data-high
description Internal Network Traffic
 match ip dscp af21 af22
!
class-map match-all out-mission-critical-mgmt
description Mission-Critical Data
 match ip dscp af31 af32
!
class-map match-all out-voice-video
description Voice/VoIP/IPT/video
 match ip dscp ef
!
policy-map out-queues
description Output queuing and scheduling policy
 class out-voice-video
  priority percent 20
  class out-mission-critical-mgmt
  bandwidth percent 40
 class out-data-high
  bandwidth percent 30
 class class-default
  fair-queue
  random-detect dscp-based
!
policy-map in-marking
 class in-voice-video
  set ip dscp ef
 class in-mission-critical-mgmt
  set ip dscp af31
```

```
 class in-data-high
  set ip dscp af21
 class class-default
  set ip dscp default
!
int loopback 0
description AS61 test
ip address 10.1.0.1 255.255.255.255
!
!
interface FastEthernet0/2
 ip address 192.168.1.1 255.255.255.252
 duplex full
 speed 100
 no shut
!
int g0/1
desc VACO LAN
ip address 10.222.70.235 255.255.255.240
max-reserved-bandwidth 100
service-policy output out-queues
service-policy input in-marking
no shut
!
int g0/2
desc UNUSED
shut
!
int g0/3
desc UNUSED
shut
!
int s1/0
desc Connection to Sprint PIP - CID
ip address 10.250.61.1 255.255.255.252
scramble
max-reserved-bandwidth 100
service-policy output out-queues
no shut
!
int s1/1
desc Connection to ATT MPLS - CID
ip address 12.2.1.1 255.255.255.252
scramble
max-reserved-bandwidth 100
service-policy output out-queues
```

```
no shut
!
!
router bgp 65061
 bgp router-id 10.1.0.1
 bgp always-compare-med
 bgp log-neighbor-changes
 bgp bestpath med missing-as-worst
 bgp bestpath as-path multipath-relax
 neighbor 10.10.10.1 remote-as 65061
 neighbor 10.250.61.2 remote-as 1802
 neighbor 12.2.1.2 remote-as 7018
 maximum-paths 2
 !
 address-family ipv4
 neighbor 10.10.10.1 activate
 neighbor 10.10.10.1 next-hop-self
 neighbor 10.10.10.1 soft-reconfiguration inbound
 neighbor 10.10.10.1 route-map no-tunnel-loopback out
 neighbor 10.250.61.2 activate
 neighbor 10.250.61.2 send-community
 neighbor 10.250.61.2 soft-reconfiguration inbound
 neighbor 10.250.61.2 route-map mpls-in in
 neighbor 10.250.61.2 route-map mpls-routes out
 neighbor 12.2.1.2 activate
 neighbor 12.2.1.2 send-community
 neighbor 12.2.1.2 soft-reconfiguration inbound
 neighbor 12.2.1.2 route-map mpls-in in
 neighbor 12.2.1.2 route-map mpls-routes out
 maximum-paths 2
 no auto-summary
 no synchronization
 network 10.220.0.0 mask 255.252.0.0
 network 10.229.0.0 mask 255.255.128.0
 network 10.229.131.128 mask 255.255.255.128
 network 199.12.113.139 mask 255.255.255.255
 exit-address-family
 !
!
router eigrp 61
network 10.0.0.0
redistribute bgp 65061 metric 1000 1 255 1 1500 route-map bgp-to-eigrp
no auto-sum
!
ip bgp-comm new-format
ip community-list 1 permit 64622:100
```

```
ip community-list 2 permit 64622:1000
ip community-list 3 deny 64622:990
ip community-list 3 permit
!
ip access-list standard to-as
 permit 10.220.0.0 0.0.15.255
 permit 10.220.16.0 0.0.7.255
 permit 10.220.24.0 0.0.3.255
 deny   10.220.0.0 0.3.255.255
 deny   10.229.0.0 0.0.255.255
 permit any
!
ip prefix-list mpls-routes seq 100 deny 0.0.0.0/0 le 32
!
ip access-list extended qos-data-high
 remark This ACL specifies all other internal VA traffic. Anything not matching
 remark this ACL will be marked with IPP (IP Precedence) 0
 remark and sent via UBR.
 remark ECSIP Reverse Proxies
 deny   ip 10.204.9.192 0.0.0.7 any
 deny   ip any 10.204.9.192 0.0.0.7
 deny   ip 10.236.0.0 0.0.31.255 any
 deny   ip 10.237.0.0 0.0.31.255 any
 deny   ip 10.238.0.0 0.0.31.255 any
 deny   ip 10.239.0.0 0.0.31.255 any
 deny   ip any 10.236.0.0 0.0.31.255
 deny   ip any 10.237.0.0 0.0.31.255
 deny   ip any 10.238.0.0 0.0.31.255
 deny   ip any 10.239.0.0 0.0.31.255
 remark Review these networks
 deny   ip host 10.254.22.59 any
 deny   ip host 10.254.22.88 any
 deny   ip any host 10.254.22.59
 deny   ip any host 10.254.22.88
 deny   ip 10.2.14.0 0.0.0.255 any
 deny   ip any 10.2.14.0 0.0.0.255
 deny   ip 10.254.16.0 0.0.0.255 any
 deny   ip 10.254.22.60 0.0.0.3 any
 deny   ip 10.254.22.64 0.0.0.15 any
 deny   ip 10.254.22.80 0.0.0.7 any
 deny   ip any 10.254.16.0 0.0.0.255
 deny   ip any 10.254.22.60 0.0.0.3
 deny   ip any 10.254.22.64 0.0.0.15
 deny   ip any 10.254.22.80 0.0.0.7
 deny   ip 10.254.32.0 0.0.3.255 any
 deny   ip any 10.254.32.0 0.0.3.255
```

```
 remark VBA client web proxies
 deny   ip host 10.205.4.21 any
 deny   ip host 10.205.4.22 any
 deny   ip host 10.205.4.23 any
 deny   ip host 10.205.4.24 any
 deny   ip any host 10.205.4.21
 deny   ip any host 10.205.4.22
 deny   ip any host 10.205.4.23
 deny   ip any host 10.205.4.24
 remark Remaining Data into Queue
 permit ip 10.0.0.0 0.255.255.255 10.0.0.0 0.255.255.255
 permit ip 10.0.0.0 0.255.255.255 152.124.0.0 0.1.255.255
 permit ip 10.0.0.0 0.255.255.255 152.126.0.0 0.0.255.255
 permit ip 152.124.0.0 0.1.255.255 10.0.0.0 0.255.255.255
 permit ip 152.126.0.0 0.0.255.255 10.0.0.0 0.255.255.255
!
ip access-list extended qos-voice-video
 remark VANTS and PolyTrauma MCU Subnet in Lees Summit.
 permit ip 10.252.243.0 0.0.0.255 any
 permit ip any 10.252.243.0 0.0.0.255
 remark Richmond AS6 Level I PolyTrauma
 permit ip host 10.49.86.142 any
 permit ip any host 10.49.86.142
 remark Tampa AS8 Level I PolyTrauma
 permit ip host 10.69.4.101 any
 permit ip any host 10.69.4.101
 remark Palo Alto AS21 Level I PolyTrauma
 permit ip host 10.168.114.241 any
 permit ip any host 10.168.114.241
 remark Minneapolis AS14 Level I PolyTrauma
 permit ip host 10.104.54.31 any
 permit ip any host 10.104.54.31
 remark Syracuse AS2 Level II PolyTrauma
 permit ip host 10.20.251.25 any
 permit ip any host 10.20.251.25
 remark West Roxbury AS1 Level II PolyTrauma
 permit ip host 10.9.111.116 any
 permit ip any host 10.9.111.116
 remark Dallas AS17 Level II PolyTrauma
 permit ip host 10.138.4.202 any
 permit ip any host 10.138.4.202
 remark Nashville AS9 Level II PolyTrauma
 permit ip host 10.78.54.89 any
 permit ip any host 10.78.54.89
 remark Houston AS16 Level II PolyTrauma
 permit ip host 10.133.204.201 any
```

```
 permit ip any host 10.133.204.201
 remark Augusta AS7 Level II PolyTrauma
 permit ip host 10.57.160.17 any
 permit ip any host 10.57.160.17
 remark Seattle AS20 Level II PolyTrauma
 permit ip host 10.161.7.71 any
 permit ip any host 10.161.7.71
 remark Denver AS19 Level II PolyTrauma
 permit ip host 10.153.99.11 any
 permit ip any host 10.153.99.11
 remark St Louis AS15 Level II PolyTrauma
 permit ip host 10.123.201.7 any
 permit ip any host 10.123.201.7
 remark Hines AS12 Level II PolyTrauma
 permit ip host 10.101.8.194 any
 permit ip any host 10.101.8.194
 remark Indianapolis AS11 Level II PolyTrauma
 permit ip host 10.89.16.44 any
 permit ip any host 10.89.16.44
 remark Cleveland AS10 Level II PolyTrauma
 permit ip host 10.83.168.13 any
 permit ip any host 10.83.168.13
 remark Washington AS5 Level II PolyTrauma
 permit ip host 10.41.246.10 any
 permit ip any host 10.41.246.10
 remark Philadelphia AS4 Level II PolyTrauma
 permit ip host 10.36.198.12 any
 permit ip any host 10.36.198.12
 remark Los Angeles AS22 Level II PolyTrauma
 permit ip host 10.180.238.1 any
 permit ip any host 10.180.238.1
 remark Tucson AS18 Level II PolyTrauma
 permit ip host 10.150.214.234 any
 permit ip any host 10.150.214.234
 remark Bronx AS3 Level II PolyTrauma
 permit ip host 10.26.249.22 any
 permit ip any host 10.26.249.22
!
ip access-list extended qos-mission-critical-mgmt
 remark Traffic for NCA WWW and gravelocator.
 permit ip host 10.228.1.2 10.238.24.0 0.0.3.255
 permit ip host 10.228.1.3 10.238.24.0 0.0.3.255
 permit ip 10.238.24.0 0.0.3.255 host 10.228.1.2
 permit ip 10.238.24.0 0.0.3.255 host 10.228.1.3
 remark VBA BDN and VACOLS
 permit ip host 10.204.10.42 any
```

```
 permit ip any host 10.204.10.42
 permit ip host 10.204.10.52 any
 permit ip any host 10.204.10.52
 permit ip host 10.222.96.30 any
 permit ip any host 10.222.96.30
 remark Network Management Subnets
 permit ip 10.254.128.0 0.0.0.255 10.0.0.0 0.255.255.255
 permit ip 10.254.129.0 0.0.0.255 10.0.0.0 0.255.255.255
 permit ip 10.254.224.0 0.0.0.255 10.0.0.0 0.255.255.255
 permit ip 10.254.28.0 0.0.0.255 10.0.0.0 0.255.255.255
 permit ip 10.254.24.0 0.0.0.255 10.0.0.0 0.255.255.255
 permit ip 10.254.40.0 0.0.3.255 10.0.0.0 0.255.255.255
 permit ip 199.12.113.0 0.0.0.255 10.254.0.0 0.0.255.255
 permit ip 10.0.0.0 0.255.255.255 10.254.128.0 0.0.0.255
 permit ip 10.0.0.0 0.255.255.255 10.254.129.0 0.0.0.255
 permit ip 10.0.0.0 0.255.255.255 10.254.224.0 0.0.0.255
 permit ip 10.0.0.0 0.255.255.255 10.254.28.0 0.0.0.255
 permit ip 10.0.0.0 0.255.255.255 10.254.24.0 0.0.0.255
 permit ip 10.0.0.0 0.255.255.255 10.254.40.0 0.0.3.255
 permit ip 10.254.0.0 0.0.255.255 199.12.113.0 0.0.0.255
!
logging source-interface Loopback0
logging 10.254.128.106
!
access-list 76 permit 10.254.128.106
access-list 76 permit 10.254.128.107
access-list 76 permit 10.254.24.107
!
access-list 91 permit 10.254.129.101
access-list 91 permit 10.254.128.101
access-list 91 permit 10.254.129.102
access-list 91 permit 10.254.128.106
access-list 91 permit 10.254.24.130
access-list 91 permit 10.254.128.4
access-list 91 permit 10.254.128.5
access-list 91 permit 10.254.128.10
access-list 91 permit 10.254.129.10
access-list 91 permit 10.254.24.107
access-list 91 permit 10.254.24.101
access-list 91 permit 10.254.24.102
access-list 91 permit 10.254.24.116
access-list 91 remark VA Management servers
!
access-list 99 permit 10.254.129.101
access-list 99 permit 10.254.128.101
access-list 99 permit 10.254.129.102
```

```
access-list 99 permit 10.254.128.106
access-list 99 permit 10.254.128.107
access-list 99 permit 10.254.24.130
access-list 99 permit 10.254.128.4
access-list 99 permit 10.254.128.5
access-list 99 permit 10.254.24.107
access-list 99 permit 10.254.24.101
access-list 99 permit 10.254.24.102
access-list 99 permit 10.254.24.66 0.0.0.60
access-list 99 permit 10.254.24.64 0.0.0.63
access-list 99 permit 10.254.27.0 0.0.0.255
access-list 99 permit 192.168.1.2
!
snmp-server community vatmp RO 91
snmp-server community vet3r@ns1st RW 76
snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server enable traps snmp authentication linkdown linkup coldstart warmstart
snmp-server enable traps envmon fan shutdown supply temperature status sonet
snmp-server enable traps config
snmp-server enable traps entity
snmp-server enable traps bgp
snmp-server enable traps rtr
snmp-server host 10.254.129.101 vatmp
snmp-server host 10.254.24.101 vatmp
!
route-map bgp-to-eigrp permit 10
 match ip address to-as
!
route-map mpls-in permit 10
 match community 1
 set metric 100
!
route-map mpls-in permit 20
 match community 2
 set metric 1000
!
route-map mpls-in permit 30
 set metric 10000
!
route-map mpls-out permit 10
 match ip address prefix-list mpls-routes
 set metric 100
 set community 64622:100
!
route-map mpls-out permit 20
```

```
 match as-path 10
 set metric 200
 set community 64622:100
!
route-map no-tunnel-loopback permit 10
match community 3
!
tacacs-server host 10.254.128.106
tacacs-server host 10.254.129.102
tacacs-server timeout 3
tacacs-server directed-request
tacacs-server key VHAWANKEY
ip tacacs source-interface loopback0
!
banner motd ^C

                        Warning

This system is for the use of authorized users only.  This system is
monitored to protect against its unauthorized use and for maintenance
purposes.  In the course of such monitoring the activities of
authorized users may be monitored and/or recorded by systems personnel.

Anyone accessing and using this system expressly consents to
system monitoring.  If monitoring reveals evidence of unauthorized
or criminal activity, systems personnel may provide that evidence to law
enforcement officials.
^C
!
line vty 0 8
access-class 99 in
password 4Va4tmp4
!
ntp logging
ntp server 10.3.31.254
ntp server 10.2.31.254
!
!enable ssh
!
cry key generate rsa   <-- needs to be done last in order to define modulus (please use 1024)
!
ip ssh time-out 60
ip ssh authentication-retries 2
!
end
```

10 APPENDIX C – DETAILED HARDWARE AND SOFTWARE BOM

10.1 Router BOM

10.2 Out of Band BOM