Network Management – Network Operations Center

Security Guidance At-a-Glance

Version 8, Release 1

Supplement of Network Infrastructure STIG, V8R1

24 March 2010

Developed by DISA for the DoD

INTRODUCTION

Network management is the process of monitoring network elements and links, configuring network elements to turn up and disable network services, the collection of performance, status, and other relevant data about each element to ensure availability, and that services are being delivered to meet or exceed service level agreements. Network management processes can be performed on site by the local network administrators and engineers or at a Network Operations Center (NOC) such as the DISN Core TNC or the Air Force Integrated Management Sites (IMS).

Whether a production network is being managed locally or from a NOC, achieving network management objectives depends on comprehensive and reliable network management solutions. These solutions provide monitoring of network behavior, availability, performance, and device configuration. Equally important is the ability to quickly detect and troubleshoot network events such as a service outage, link down, node down, and high utilization of both network elements and circuits.

The intent of this document is to discuss the best network management practices that should be implemented and to provide easy to follow guidance to securely manage networks. It is a supplement to the Network Infrastructure STIG and will include current network management requirements from the STIG as well as new requirements that have been integrated into the Network Checklist. The new requirements found in this document that have been integrated are as follows:

NET0435	NET0817	NET0990	NET0997	NET1004	NET1617
NET0436	NET0819	NET0991	NET0998	NET1005	NET1731
NET0437	NET0985	NET0992	NET0999	NET1006	NET1732
NET0438	NET0986	NET0993	NET1000	NET1007	NET1733
NET0809	NET0987	NET0994	NET1001	NET1008	NET1734
NET0814	NET0988	NET0995	NET1002	NET1615	NET1807
NET0815	NET0989	NET0996	NET1003	NET1616	NET1808
NET0816					

A large portion of the guidance found in Section 1 will be dependent on network topology as well as the specific network management deployment based on various out-of-band (OOB) and in-band management paradigms. Hence, the flowchart below provides and easy to follow path that will enable the reader to determine what must be implemented.

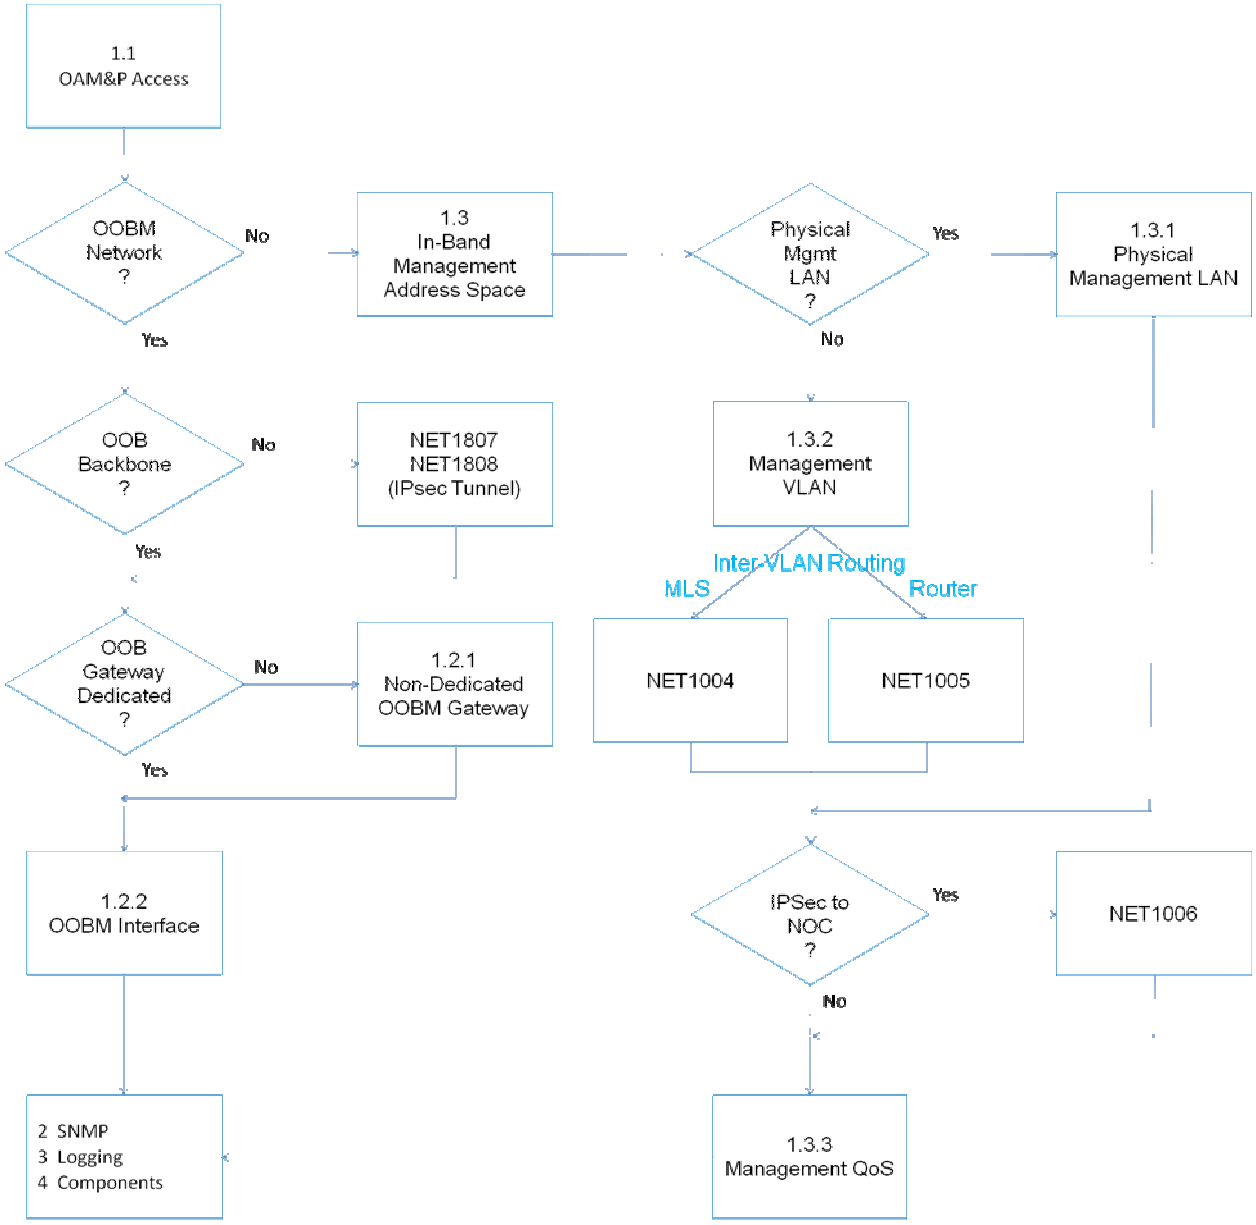
1.1
OAM&P Access
OOBM
Network
?
No
Yes
1.3
In-Band
Management
Address Space
Physical
Mgmt
LAN
?
Yes
No
1.3.1
Physical
Management LAN
OOB
Backbone
?
No
Yes
NET1807
NET1808
(IPsec Tunnel)
1.3.2
Management
VLAN
Inter-VLAN Routing
MLS
Router
OOB
Gateway
Dedicated
?
No
Yes
1.2.1
Non-Dedicated
OOBM Gateway
NET1004
NET1005
1.2.2
OOBM Interface
IPSec to
NOC
?
Yes
No
NET1006
2 SNMP
3 Logging
4 Components
1.3.3
Management QoS

1 MANAGEMENT NETWORK

Management systems provide the network operator the facility to manage the network and all of its components. They are both the platforms and applications that interact with the managed network elements to provide the NOC a framework to facilitate operation, administration, maintenance, and provisioning (OAM&P) tasks. OAM&P is a group of management functions that enable system or network fault indication, and diagnostics, performance monitoring, security management, configuration management, and service provisioning. Management systems and managed network elements need to be interconnected. The facility that provides this connection is referred to as the management network.

To be managed, a network element provides a management interface through which a management system can communicate. Hence, the management system is the reason for the management network to exist. The management network is comprised of network management workstations, authentication servers, syslog servers, communications servers, Operations Support System (OSS), and a network for transporting management traffic. While the Network Infrastructure STIG provides guidance for securing a network, including the management network, this section will discuss the connectivity models used to access the network being managed as well as all of the management network components, the vulnerabilities that they introduce, and what security measurements must be taken to mitigate these risks.

1.1 Network Element Access for OAM&P

To provide management access, network elements support direct serial connections, out-of-band connections, and in-band connections. Either in-band or out-of-band connections are used to transport network management messages between the managed network elements and management systems used for providing OAM&P functions. In either case, the same services such as telnet, secure shell, http, and SSL are used to access a managed network element. Out-of-band and in-band management implementations are reviewed in Sections 1.2 and 1.3 respectively.

The direct serial interface is typically referred to as the craft port or console port. There may also be an auxiliary port. This interface is intended to be an access port through which code downloads and local monitoring and control can take place. The auxiliary port, consol port, as well as any slow-speed async serial port with an analog modem connected to it provides the capability for direct dial-up administrative access. If dial-up capability is provided, a secured modem and connection must be used as specified in Section 4.2.

1.2 Out-of-Band Management Network

The Out-of-Band Management (OOBM) network is an IP network used exclusively for the transport of OAM&P data from the network being managed to the OSS components located at the NOC. Its design provides connectivity to each managed network element enabling network management traffic to flow between the managed network elements and the NOC. This allows

the use of paths separate from those used by the network being managed. The NOC could be located locally or remotely at a single or multiple sites all connecting to the OOBM network.

OOBM networks isolate network users from communication channels that are dedicated to network management. As shown in Figure 1, all managed devices are connected to the OOBM access switch via the managed elements' OOBM interface. This OOBM access switch must be used only for OOBM and therefore provide access to only the OOBM router, the OOBM communications server, and the OOBM interfaces of the managed network elements. The OOBM access switch, the OOBM gateway router and comm server interface connecting to the access switch, and all of the managed network elements' OOBM interfaces are essentially the OOBM remote site LAN that has its own subnet and hence its own broadcast domain.

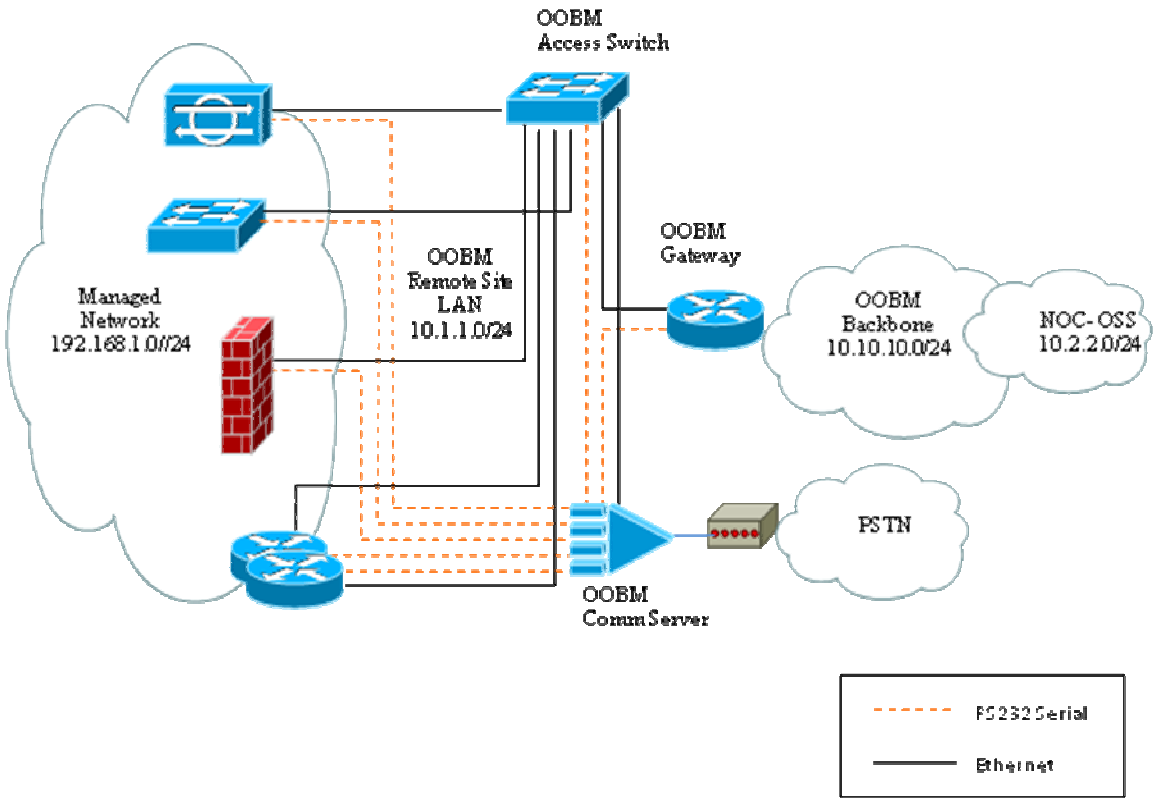

Figure 1 OOBM Access to Managed Network

The OOBM router is the gateway between the network elements being managed and the OOBM backbone. Using dedicated paths, the OOBM backbone connects the OOBM gateway routers located at the premise of the managed networks and at the NOC. Dedicated links can be deployed using provisioned circuits (ATM, Frame Relay, SONET, T-carrier, and others or VPN technologies such as subscribing to MPLS Layer 2 and Layer 3 VPN services) or implementing a

secured path with gateway-to-gateway IPsec tunnels. Figure 2 illustrates a topology with the NOC and the OOBM remote site LAN securely connected via IPsec tunnels between the OOBM gateway routers.

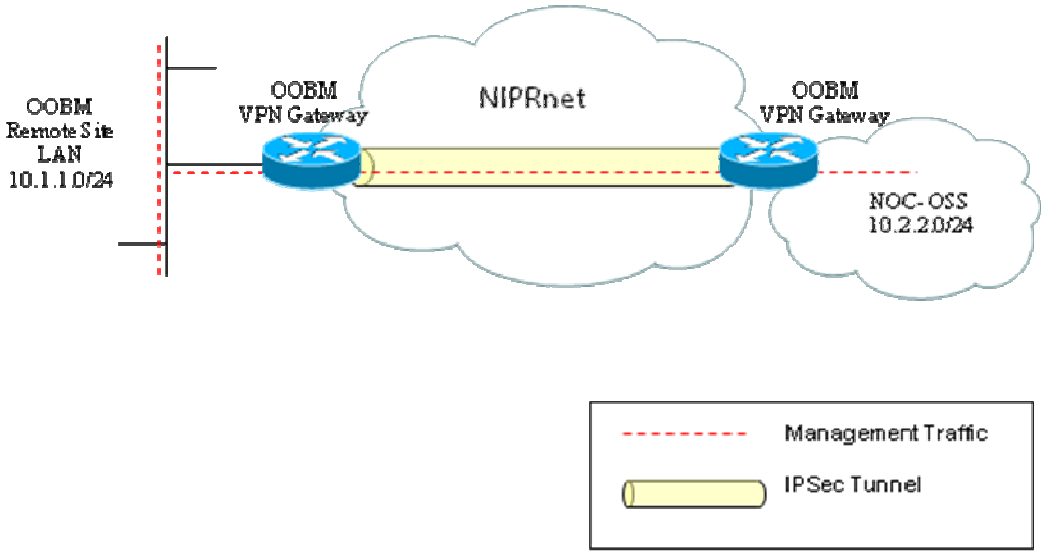

Figure 2 Remote Site OOBM Connectivity via IPsec

1.2.1 Non-Dedicated OOBM Gateway Routers

If the gateway router is not a device dedicated for the OOBM network (i.e. may be the managed network's premise router), several safeguards must be implemented for traffic containment and separation. Management traffic must not leak into the managed network, and traffic from the managed network must not leak into the management network. Since the managed network and the management network are separate routing domains as depicted in Figure 3, separate IGP routing instances must be configured on the router—one for the managed network and one for the OOBM network. In addition, this router must be configured to ensure that control plane traffic is not redistributed between the two routing domains.

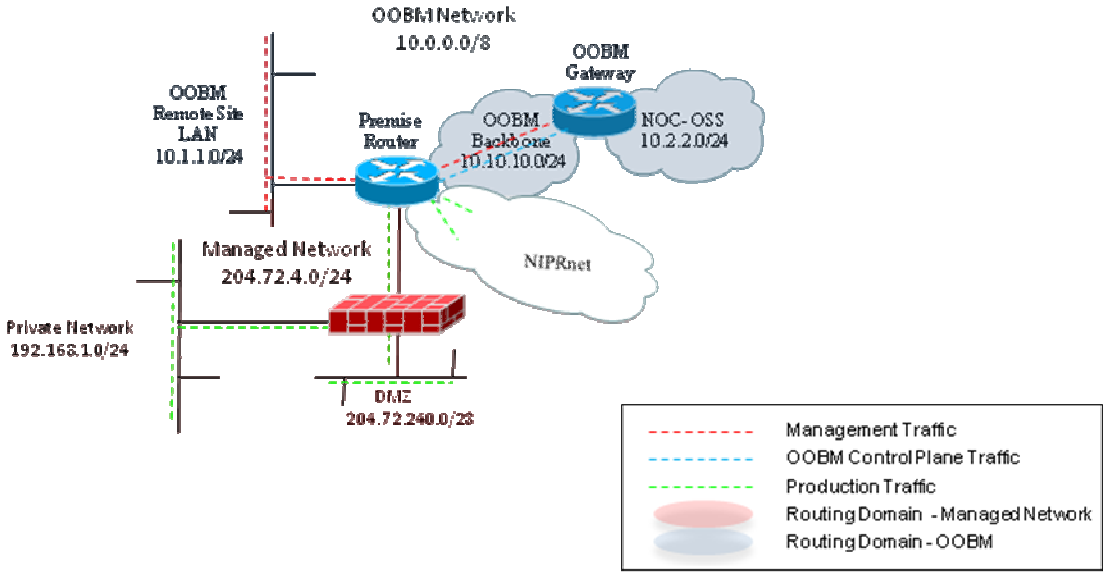

Figure 3 Non-Dedicated OOBM Gateway – Dedicated OOB Backbone

For OOBM deployments lacking dedicated OOB backbone links as previously discussed, secured paths can be deployed using IPsec tunnels between the gateways. If static routing is utilized, implementing IPsec tunnel between the non-dedicated OOBM gateway and the OOBM gateway at the NOC to transport management traffic is the only requirement. However since static routes do not scale well, dynamic routing may be required. Hence, control plane traffic must be able to traverse the secured path. This deployment is implemented by establishing a GRE tunnel between the two gateways. IGP routing protocol adjacencies will form over the GRE tunnel end points. The GRE tunnel will be encrypted by IPsec to provide privacy for the control plane payload as illustrated in Figure 4.

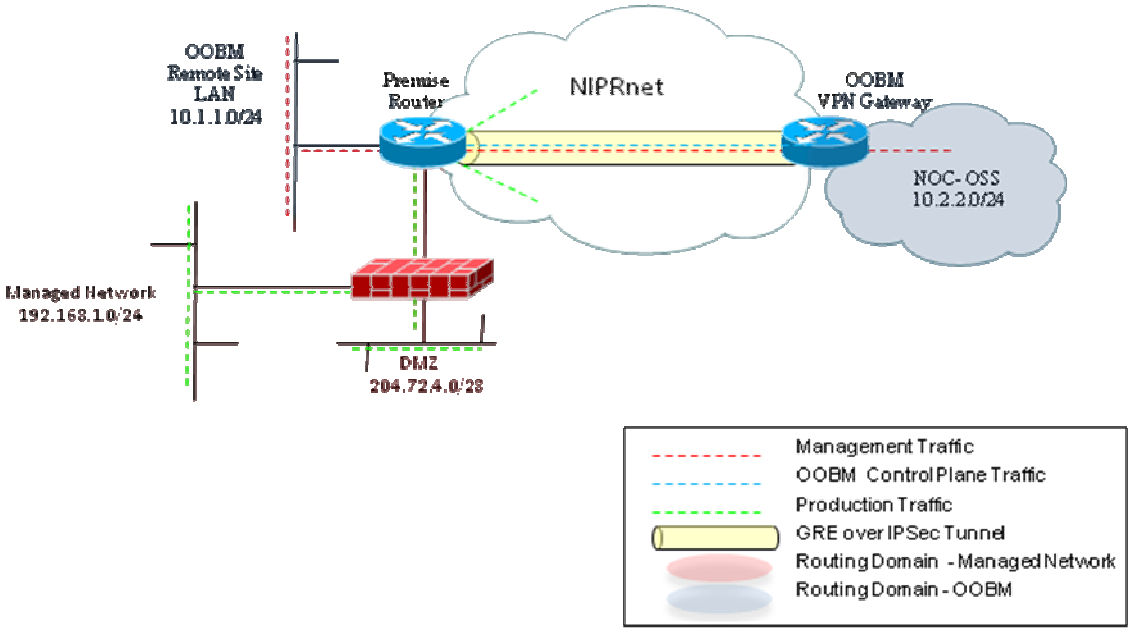

Figure 4 Non-Dedicated OOBM Gateway - Connectivity via GRE/IPsec

1.2.2 OOBM Interface

The OOBM access switch will connect to the management interface of the managed network elements. The management interface can be a true OOBM interface or a standard interface functioning as the management interface. In either case, the management interface of the managed network element will be directly connected to the OOBM network.

An OOBM interface does not forward transit traffic; thereby, providing complete separation of production and management traffic. Since all management traffic is immediately forwarded into the management network, it is not exposed to possible tampering. The separation also ensures that congestion or failures in the managed network do not affect the management of the device. If the device does not have an OOBM port, the interface functioning as the management interface must be configured so that management traffic does not leak into the managed network and that production traffic does not leak into the management network.

1.3 In-Band Management Network

The in-band management paradigm exists when the management traffic takes the same data path as non-management or production traffic, thereby using the same physical or logical interface. Management plane traffic share the same path as the control plane and forwarding plane. Henceforth, network management traffic is intermixed with user traffic using the same interfaces of the network elements being managed.

Applications used to access the managed device utilize TCP as the transport protocol. The TCP keepalive feature will periodically verify that the remote node of the management session (i.e., the network management station) is still reachable. In the event the remote node has abnormally terminated or an upstream link from the managed device is down, the management session will be terminated; thereby, freeing device resources and eliminating any possibility of an unauthorized user being orphaned to an open idle session of the managed device.

Unlike out-of-band implementation, the configured IP address of the interfaces used to access the devices belong to address space of the managed network. Using a loopback address as the source address provides security, scalability, and manageability of all routers and switches. It is easier to construct appropriate ingress filters for management plane traffic destined to the network management subnet since the source addresses will be from the range used for loopback interfaces as to the larger range used for physical interfaces. Log information recorded by authentication and syslog servers will record the router's loopback address instead of the numerous physical interface addresses. Messages sent to the following servers must also use the loopback address as the source address: Syslog, TACACS+, RADIUS, NTP, SNMP, NetFlow Collector, TFTP, and core dumps to FTP servers.

1.3.1 Physical Management LAN Segment

As illustrated in Figure 5, the management network must still have its own subnet in order to enforce control and access boundaries provided by Layer 3 network nodes such as routers and firewalls. Management traffic between the managed network elements and the management network is routed via the same links and nodes as that used for production or operational traffic. Safeguards must be implemented to ensure that the management traffic does not leak past the managed network's premise equipment.

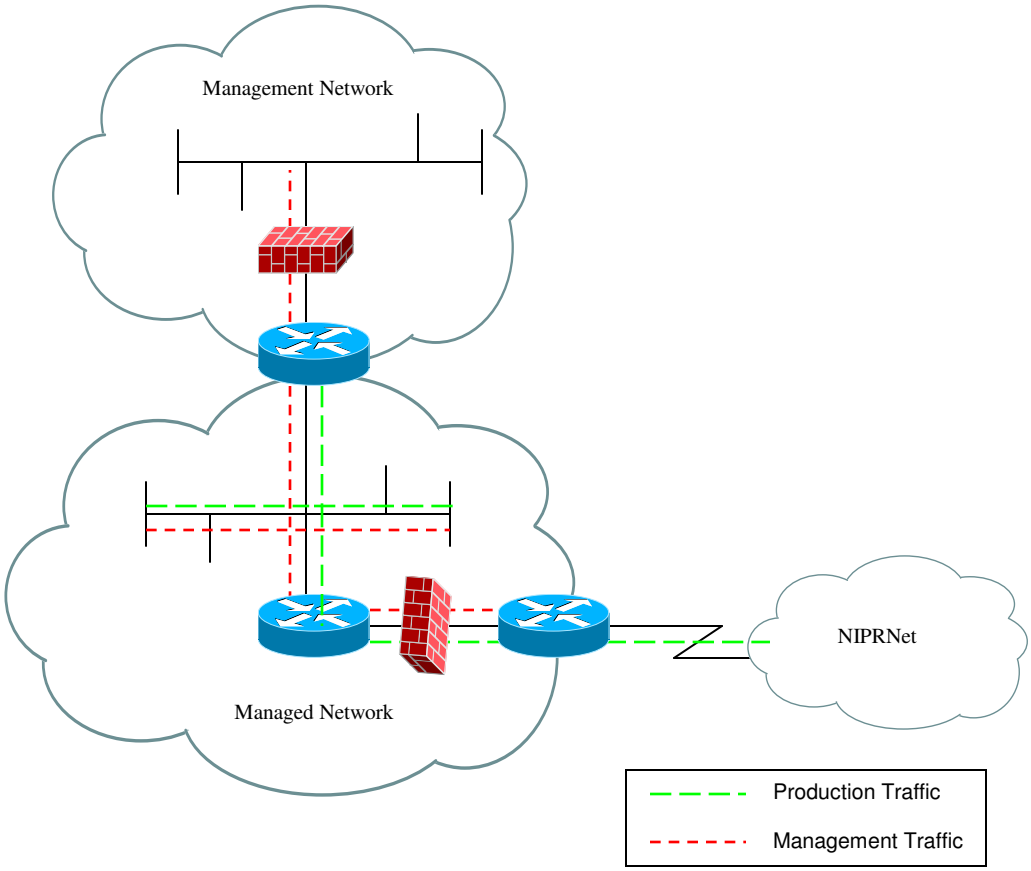

Figure 5 Management Network Separation

1.3.2 Management VLAN

If the management systems reside within the same layer 2 switching domain as the managed network elements, then separate VLANs will be deployed to provide separation at that level. In this case, the management network still has its own subnet while at the same time it is defined as a unique VLAN. As illustrated in Figure 6, inter-VLAN routing or the routing of traffic between nodes residing in different subnets requires a router or multi-layer switch (MLS). Access control lists must be used to enforce the boundaries between the management network and the network being managed. All physical and virtual (i.e. MLS SVI) routed interfaces must be configured with ACLs to prevent the leaking of unauthorized traffic from one network to the other.

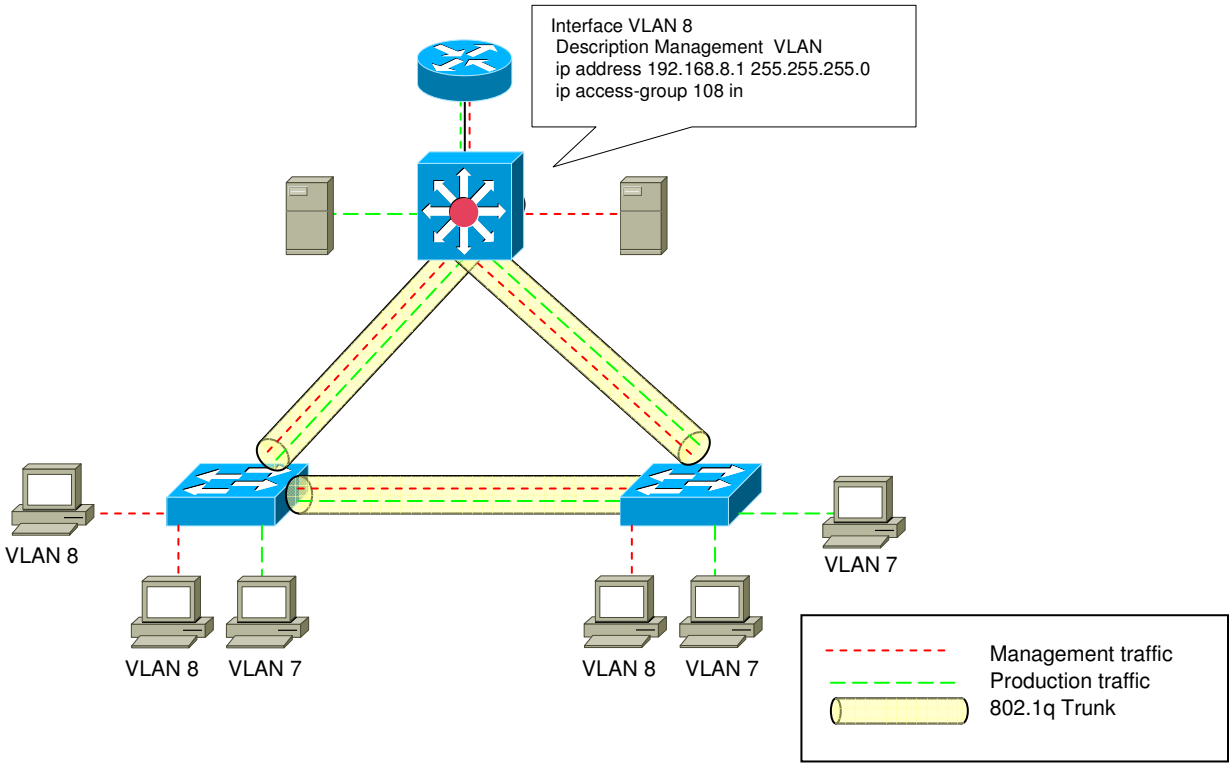

Figure 6 Management VLAN Separation

1.3.3 NOC Connectivity

Similar to the OOBM model, when the production network is managed in-band, the management network could also be housed at a NOC that is located locally or remotely at a single or multiple interconnected sites. NOC interconnectivity as well as connectivity between the NOC and the managed networks' premise routers would be enabled using either provisioned circuits or VPN technologies such as IPsec tunnels or MPLS VPN services. The topology shown in Figure 7 depicts all management traffic between the NOC sites and the managed network encapsulated within IPsec tunnels traversing the NIPRNet.

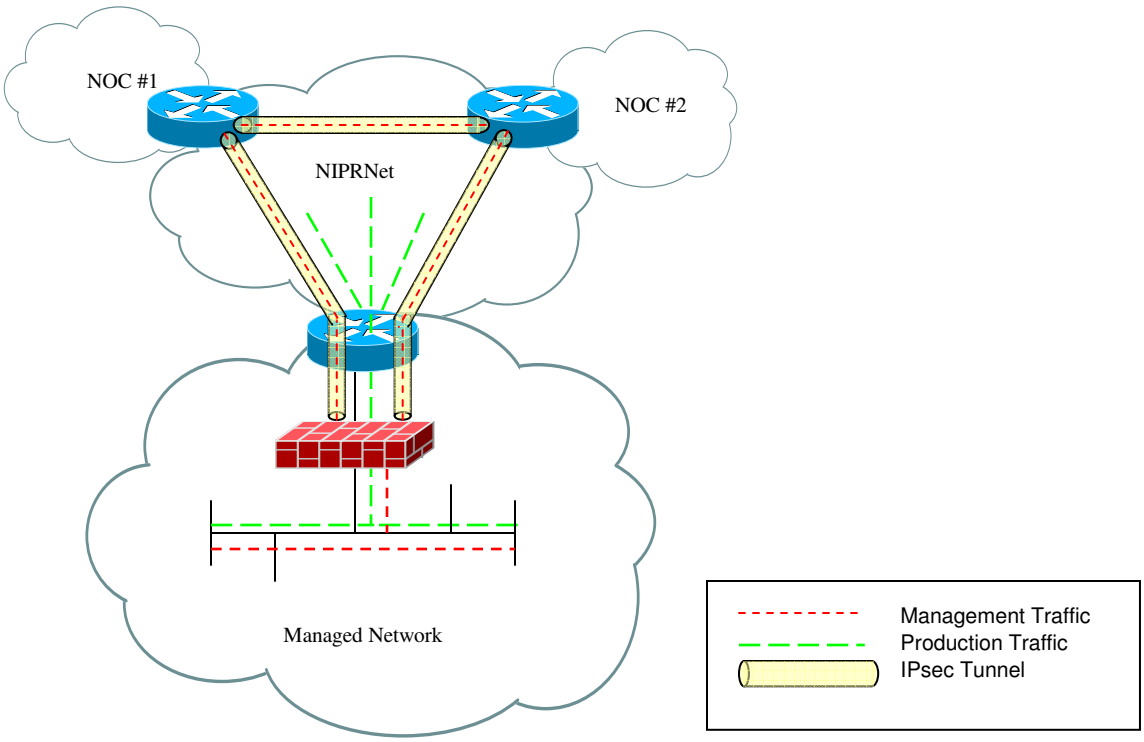

Figure 7 Traffic Separation

1.3.4 Management Traffic QoS

When network congestion occurs, all traffic has an equal chance of being dropped. Prioritization of network management traffic must be implemented to ensure that even during periods of severe network congestion, the network can be managed and monitored. Quality of Service (QoS) provisioning categorizes network traffic, prioritizes it according to its relative importance, and provides priority treatment through congestion avoidance techniques. Implementing QoS within the network makes network performance more predictable and bandwidth utilization more effective. Most important, since the same bandwidth is being used to manage the network, it provides some assurance that there will be bandwidth available to troubleshoot outages and restore availability when needed.

When management traffic must traverse several nodes to reach the management network, management traffic should be classified and marked at the nearest upstream MLS or router. In addition, all core routers within the managed network must be configured to provide preferred treatment based on the QoS markings. This will ensure that management traffic receives preferred treatment (per-hop behavior) at each forwarding device along the path to the management network.

2 SIMPLE NETWORK MANAGEMENT PROTOCOL

SNMP version 3 (SNMPv3) provides secure exchanges of management data between network devices and network management systems. The encryption and authentication features in SNMPv3 ensure high security in transporting packets to a management console. SNMPv3 employs the User-based Security Model (USM) to provide cryptographic services. The USM currently uses either MD5 or SHA message digests to ensure message authenticity and integrity, and AES encryption to ensure message privacy. These features are used to provide three distinct levels of security: 1)no authentication with no privacy, 2)authentication with no privacy, and 3) authentication with privacy. Data can be collected securely from SNMP devices without fear of the data being tampered with or corrupted. SNMP Set command packets that change a router's configuration can be encrypted to prevent its contents from being exposed on the network.

3 LOGGING

Logging is a key component of any security architecture and is a critical part of network element security. It is essential security personnel know what is being done, what was attempted, and by whom in order to compile an accurate risk assessment. It is also imperative that all configuration changes to network elements are logged on a per-session and per-user basis. Maintaining an audit trail of system activity logs can help identify configuration errors, understand past intrusions, troubleshoot service disruptions, and react to probes and scans of the network.

4 NETWORK MANAGEMENT AUXILIARY COMPONENTS

The network management auxiliary components are used to provide capabilities to enable both management and security functionality for the managed network. These components are being secured as a result of the IA requirements that have been defined based on the topology—that is, whether they are residing within a dedicated OOB network infrastructure or connected to an in-band network. Nevertheless, since they do have sessions with elements in the managed network that could be compromised, there are additional IA measures that must be followed to reduce the risk of these components from also being compromised.

4.1 Syslog Server

A syslog server provides the network administrator the ability to configure all of the communication devices on a network to send log messages to a centralized host for review, correlation, reporting, and storage. This implementation provides for easier management of network events and is an effective facility for monitoring and the automatic generation of alert notification. The repository of messages facilitate troubleshooting when problems are encountered and can assist in performing root cause analysis. Syslog files can also be parsed in real time to identify suspicious behavior or be archived for review at a later time for research and analysis.

4.2 Communications Servers

A communications server (aka terminal server) can be used to provide interconnectivity between all managed network elements and the OOBM gateway router for administrative access to the device's console port. In the event the OOBM network is not able to provide connectivity due to an outage, the communications server can provide a dial-up PPP connection to access a network element. The auxiliary port, consol port, as well as any slow-speed async serial port with an analog modem connected to the managed device also provides the capability for direct dial-up administrative access for infrastructures that do not have a communications server for management access.

4.3 Authentication Server

Using standardized authentication protocols such as RADIUS, TACACS+, and Kerberos, an authentication server provides centralized and robust authentication services for the management of network components. An authentication server is very scalable as it supports many user accounts and authentication sessions with the network components. It allows for the construction of template profiles or groups that are given authorization for specific tasks and access to specific resources. Users are then given an account that has been configured in the authentication server and as been assigned to a group.

4.4 NTP Client & Server

NTP provides an efficient and scalable method for network elements to synchronize to an accurate time source refered to as the reference clock or stratum-0 server. The reference clock synchronizes to the Coordinated Universal Time (UTC) derived from a set of atomic clocks

using GPS, CDMA, or other time signals such as Irig-B, WWV, and DCF77. A stratum-0 server cannot be connected to a network. Instead, it is directly connected to an IP network-enabled host which then operates as a stratum-1 NTP server. NTP time distribution is based on a loop-free topology with the stratum-1 server as the root of the tree—this includes NTP servers with built-in stratum-0 components. NTP updates can be sent as unicast, multicast, or broadcast. The later methods enables a server to synchronize multiple hosts in an unsolicited mode.

Every router that receives and accepts time from a stratum-n server automatically becomes a stratum-n^{+1} server. An NTP node will not accept time updates from an NTP server at a higher stratum; thereby enforcing a tree-level hierarchy of client-server relationships and preventing time syncronization loops. Two NTP servers at the same stratum may form peer relationships and thus provide time to each other. An NTP topology should be designed to easily scale by creating a stratum hierarchy of servers to accommodate the workload. The width (number of servers at same stratum level) and depth (number of stratum levels or tiers) of the hierarchy is dependent on the number of NTP clients as well as the amount of redundancy that is required.

It is vital for network operations and security management to prevent unauthorized time sources from altering or interferring with time synchronization within the managed network. Depending on vendor implementation, NTP-eneabled network devices will accept any node as their peer. Hence, a rouge device could pose as an NTP peer and begin sending false time to a router. To ensure that manged devices do not recive time from imposters, they must be configured with access control lists to restrict by IP address which servers and peers a device will accept NTP update messages from.

To launch an attack on the NTP infrastructure, a hacker could inject time that would be accepted by NTP clients by spoofing the IP address of a valid NTP server. To mitigate this risk, the time messages must be authenticated by the client before accptimg them as a time source. The NTP authentication model is opposite of the typical client-server authentication model. NTP authentication allows NTP cleints to authenticate their servers and peers. It's not used to authenticate NTP clients because NTP servers don't care about the authenticity of their clients, as they never accept any time from them.

Insuring that there are always NTP servers available to provide time is critical. It is imperative that all single points of failure for the NTP infrastructure are eliminated. Network nodes syncronizing to UTC is crucial for network operations as well as for security management (forensics, auditing, certificate expirations, time-based access control, etc). Compromising an NTP server opens the door to more sophisticated attacks that include NTP poisoning, replay attacks, and denial of service.

To provide security through separation and isolation, the NTP server should only be connected to the management network. This enables the NTP server to provide time to the managed devices using a secured as well as a preferred path. If the NTP server is not an appliance, it is critical that the system is secured by maintaining compliance with the appropriate OS STIG as well as implementing an HIDS.

4.5 SNMP Manager

The SNMP manager provides the interface between the network management personnel and the managed network. On the other hand, the SNMP agent provides the interface between the manager and the device being managed. The manager is the collector of alarm information via SNMP traps as well as statistical and historical management information retrieved by polling the agents within the managed network. This information is vital for real time monitoring and alarm management as well as for strategic planning and performance management. In addition to the SNMP safeguards outlined in section 2, IA measures must be implemented to mitigate the risk of the SNMP manager being compromised. To provide security through separation and isolation, the SNMP manager should only be connected to the management network. This enables the SNMP manager to provide management services to the managed devices using a secured as well as a preferred path.

5 LOGISTICS: IMAGE AND CONFIGURATION STORAGE

It is important to keep the running configuration and the startup configuration synchronized, so that if there is a power failure or some other problem that forces the device to restart, the managed router or switch will load the correct configuration. If there is a need to rollback to an older configurations, they should be stored offline on a FTP or TFTP server.

Images installed on the devices flash memory can become corrupt. Hence, it is imperative to retain a copy of the current production images on some form of offline media or file server. Both prior and new image versions should also be kept in the event regression occurs or for planned upgrade migrations respectively. With the image and configuration files stored offline, the files must be transferred to and from the switch or router in a secure method. FTP is preferred over TFTP, provided that the FTP server requires client authentication. Following are some alternative approaches that are actually more secure than using FTP:

- Copies of the device configuration can be archived on the devices' flash or hardrive if the media is available.

- If the router or switch is equipped with a PCMCIA flash memory card, images as well as configurations can be copied unto the card and stored offline for backup purposes.

- Copy and paste output of a displayed configuration while in a SSH session or HyperTerminal console connection. The file can then be saved onto a floppy disk and stored in a secure location.

- Use Secure Copy Protocol (SCP) which requires that authentication, authorization, and accounting (AAA) be configured in order for the router or switch to determine whether the user has the correct privilege level.