Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

Interested in learning more?

Check out the list of upcoming events offering
"Auditing Systems, Applications, and the Cloud (Audit 507)"
at http://www.giac.org/registration/gsna

SANS GSNA Certification
GSNA – Auditing Networks, Perimeters, and Systems (v2.1)

Auditing hp OpenView Network Node Manager
An Auditors Perspective

Practical Assignment
October 2003

Submitted by: Josh O'Mealey

Table of Contents

3

Assignment 1 – Research in Audit, Measurement Practice, and Control

1.1 – Abstract

    The purpose of this paper is to discuss the auditing steps and procedures used when auditing an HP OpenView Network Node Manager (NNM) system. The paper was written from the perspective of a third-party auditor, but current NNM administrators will benefit from the best-practice and hardening steps found throughout the document. An auditing checklist concerning security testing and best-practice hardening settings, for an NNM system and server, is discussed in detail.

    NNM systems and servers are at great risk in a live environment. There are many vulnerabilities that an attacker could exploit on an NNM system or server. Enterprises treat NNM systems as mission-critical systems. These systems should be tested persistently using the auditing steps, which are discussed in detail throughout the audit checklist, to ensure that the NNM system and server are protected. The scope of this paper is to demonstrate the auditing steps needed to successfully audit an NNM system and server.

    This paper will also give the reader a quick overview of NNM functions and the Simple Network Management Protocol (SNMP). SNMP concepts and the protocol's enterprise-wide security threats will be discussed throughout this paper. NNM is tightly integrated with the SNMP and depends on this protocol to realize its full functionality. Testing results and recommendations will be presented in a management summary after the audit has been carried out on a live NNM system and server.

1.2 – Identify the System to be Audited

    Hewlett Packard OpenView Network Node Manager v6.4 is a high-end network management and monitoring solution for enterprises and Internet Service Provider environments. NNM utilizes Simple Network Management Protocol (SNMP) to automatically discover, manage, and monitor all SNMP-enabled devices on a Local Area Network (LAN) and/or Wide Area Network (WAN). NNM is also able to monitor any device on a network that is running Internet Protocol (IP) or Internet Packet Exchange (IPX).

    Most enterprises that use NNM will use it in one of three ways: as a network device monitoring solution, as a network device management solution, or a combination of the two. The NNM system that will be audited is currently being used as the main network device monitoring tool for an enterprise LAN/WAN environment. The NNM administrator has also integrated an Urgent Messaging System (UMS) into the NNM server.

4

Several teams within this organization rely on NNM to accurately relay network device status messages through the UMS, which are then sent out to network administrators. For NNM to be able to fully monitor the equipment on the network, it is treated as a trusted machine that has access to poll equipment for status reports and topology information anywhere on the network.

This level of trust requires the NNM server to be allowed access through network level Access Control Lists (ACL's) that are enforced on network routers, so that the NNM server is able to see everything on the network. With this level of trust comes a great deal of responsibility for the NNM administrator to have a very secure NNM system and server that is running the trusted application.

Model:	HP OpenView Network Node Manager
Version:	6.4
Version Number:	B.06.41
Operating System:	Microsoft Windows 2000 Server
Operating System Number:	Build 2095 / Service Pack 3
Role:	Network Monitoring Tool
Description:	The NNM application monitors all SNMP-enabled devices that are important to the network. All routers, switches, servers, and print server appliances are monitored for interface and device up and down status.

The NNM server that will be audited is a Microsoft Windows 2000 Server with an Intel Pentium III CPU running at 1.4 Ghz and 1,310 MB of physical RAM. It also has a 10/100 Mbps Network Interface Card (NIC) and a 100 Mbps full-duplex LAN connection. The system has Virtual Memory settings of 3,836 MB.

1.2 - Evaluate the Risk to the System

The risk associated with an NNM system is directly related to how the product is being used. If it is being used as a network management solution, then the SNMP Read/Write (RW) string is present on the network devices, which drastically increases the risk to the NNM system. If it is being used as a network monitoring solution, then only the SNMP Read Only (RO) string is being used on the devices, thus the ability for an attacker to cause damage after gaining control of an NNM system is drastically reduced.

A good attacker knows that if access can be gained to an NNM system that is managing network devices much more damage can be done, compared to an NNM system strictly being used for monitoring. If an attacker was able to gain control of a system that could perform SNMP RW commands on devices all over the network, the core architecture of the network could essentially be taken over and the system administrators locked out from their devices.

5

This is not to say that a knowledgeable attacker is not still interested in gaining control of an NNM system that is being used purely for network monitoring. If an attacker were able to complete this task, they would have an advanced understanding of the entire mission-critical network. The attacker would no longer need to go through the usual preliminary phase of attack, which is purely dedicated to reconnaissance. This is by far the longest and most agonizing part of a skilled attack, and every attacker would like to scrape some time off of this phase and move on to attacking key systems.

The NNM system that is being audited is strictly being used as a monitoring tool for the network administrators. The administrators are concerned from a security standpoint with having a system that can control all of the devices on a network from one central point. The NNM system is primarily being used to monitor the status of the mission-critical devices on a network. The NNM system must always be up and functioning well to give the network administrators confidence that all LAN and WAN site Service Level Agreements (SLAs) are being fulfilled with high uptime percentages.

NNM relies on SNMP traffic from monitored devices for the enterprise to realize the full functionality of the product. SNMP traffic passes over Transmission Control Protocol (TCP) ports 161 and 162 and User Datagram Protocol (UDP) ports 161 and 162. Since the introduction of the first version of the SNMP protocol, all TCP and UDP packets have been sent in cleartext across monitored networks. SNMPv2c was introduced with new functionality, but still lacked data encryption methods for packet transmission security.

SNMPv3 has recently been released, and it introduces data encryption capabilities to the protocol. Many new SNMP-enabled devices are being shipped with the SNMPv3 protocol installed, but NNM does not yet support SNMPv3 natively. However, SNMP Research International (www.snmp.com) has released an add-on pack for NNM called SNMPv3 Security Pack (http://www.snmp.com/products/snmpsecpack.html). This will install locally on an NNM server and will integrate seamlessly into any NNM system running versions 4.1 and higher.

The SNMPv3 protocol was engineered to use data encryption algorithms during all SNMP data transmissions. Before SNMPv3 was introduced SNMP packets were easily captured and read, since the SNMP trap was being sent in cleartext. Intercepted SNMPv2c packets sent to and from monitored SNMP-enabled devices could give an attacker valuable information about the system that sent the packet and the network with which it was connected.

The SNMPv3 Security Pack will not be discussed in the audit checklist in Assignment 2, because this would be outside the scope of this paper, which is intended to audit a default NNM installation. NNM administrators however, should

look into purchasing this piece of software and integrating it into the enterprise's current NNM system. This will add a great deal of security to the monitored network concerning secure data transmissions from mission-critical devices on the monitored network.

What Can Go Wrong	*How Likely is it to Happen*	*What are the Consequences*
Remote Control	An unauthorized user is able to install an NNM remote management client, and control the NNM system. <u>High</u>	The system can become unreliable and unstable if manipulated incorrectly, whether intentionally or unintentionally.
Web Server Hack	NNM has to have a web server running, and Microsoft's Internet Information Service (IIS) is installed by default. <u>High</u>	The attacker could gain full control of the server and destroy the NNM installation.
Internal Access	An unauthorized user is able to open the web maps and then know the entire enterprise topology. <u>High</u>	The user could use this for reconnaissance to plan an attack in the future.
Denial of Service Attack	Loss of accurate node status throughout the enterprise. <u>Medium</u>	If mission-critical devices can no longer server requests and the appropriate system administrators are not alerted to it quickly, then network downtime could be seen.
System Administrator Error	Vulnerability patches could corrupt the installation. <u>Low</u>	There could be a lapse in network status notification if the patch has to be removed after the reboot.
SNMP Packet Capture	SNMP traps could be intercepted during a malicious packet capture session. <u>Low</u>	This would be difficult to do, but critical system information and SNMP configurations of a device could be captured and read in the clear.

1.3 - What is the current State of Practice?

After much searching on the Internet, I found that very little if any has been written about an NNM system from a best-practice or auditing checklist perspective. The only paper that I could find that came anywhere close to the same kind of research that I was looking for, was that of Rich Antonick's GCUX paper on a secure installation of NNM on Unix (http://www.giac.org/practical/Rich_Antonick_GCUX.doc).

- Resources Used in Research:

7

- o HP - OpenView Managing Your Network for NNM
- o HP OpenView Scalability and Distribution for NNM
- o CERT – www.cert.org
- o HP – openview.hp.com
- o SANS – www.sans.org
- o SNMP - www.snmp.com
- o Google – www.google.com

CERT (www.cert.org) had several vulnerabilities posted, but all were concerning NNM versions prior to 6.4. It seemed as though HP was doing a better job with their more recent releases of NNM, by correcting flaws from their previous releases. On average HP releases a vulnerability patch about once a month.

There are some good access control recommendations from HP in the 900 pages of documentation that is sent along with each major version release. The problem for most NNM administrators is that the security benefits are hinted at, and spread throughout the chapters. I have had to read every page in the past to feel like everything was being done to take care of the security of an NNM system that I was maintaining.

The audit checklist that follows contains information that has been pulled from several sources, but used in a different way than what may have been intended in the original printing. Several steps are advanced NNM configurations that aren't really discussed as security hardening methods within an NNM system, but I feel they are presented with that perspective below. Also, many of the steps will come from my own personal experience in securing a current NNM implementation in a large enterprise environment, primarily because of the difficulty in finding this kind of information anywhere else.

Assignment 2 – Create an Audit Checklist

2.1 - Network SNMP Checks

The section composed of the audit checklist will assume that the auditor has asked the NNM administrator to produce a list of authenticated NNM users. If the NNM administrator is the only user on the system, then it should be so noted in the document. If there are other NNM users, their domain login names and access level privileges should be listed.

Step 1: Check for any unknown SNMP-enabled devices on the network
Reference: Personal Experience
Control Objective: This step checks to see if there are any SNMP-enabled devices out on the network that are configured with *Public* as the RO string. This will need to

8

be done at the beginning of the audit, because of time/discovery constraints.

Risk: If any SNMP-enabled devices have been placed on the network with a default RO string, those devices need to be found and reconfigured as quickly as possible. If *Public* is being used as the RO string on a device it needs to be changed or the SNMP service needs to be disabled, because an attacker could easily access specific information about the configuration of the device or certain information of the network itself.

Compliance: The SNMP configuration of a device can quickly be checked locally or across the network, which is why it is so easy for an attacker to do the same. The auditor will verify whether the device is in compliance with the enterprise's policy, which would discuss not leaving the default SNMP configuration on a device. If NNM picks up the device using the default RO string, the auditor would then know that it was not in compliance.

Testing: To find out if there are any devices on the network that are accessible using *Public* as the RO string, the auditor should open the file *$NNMserver\conf\netmon.cmstr*. This file will define any RO or RW strings, other than the global default, that are acceptable to the NNM system for automatic device discovery purposes. *Public* should be placed at the bottom of the list of acceptable RO strings ("*Public*" ::::). The command *ovstatus –v netmon* should be run on the command line of the NNM system to see what a normal output looks like on the current system configuration. The auditor needs to know the status of the polling and DNS lookups (if the system is already behind, the system may not be able to accelerate IP discovery). The auditor and NNM administrator should then slowly decrement the discovery interval value that is set for the NNM system to ensure that any new devices will be found quickly for testing purposes, but without overloading the NNM system or network with IP discovery traffic. The NNM system's hardware capabilities and network bandwidth will directly affect how fast it can locate new devices that have been placed out on the network recently. If the NNM system's discovery interval is decremented too quickly, the system will get behind in its other activities because it is trying to touch every corner of the LAN within the specified time period. Polling and status updates on devices will become inaccurate, because the NNM system is using too many of its resources to find new devices and not enough resources trying to keep track of monitored devices' status accurately. The IP Discovery interval can be changed by pulling up the RW map, and then navigating to Options in the File menu, Network Polling Configuration, and then select the IP Discovery tab within the subsequent window. The IP Discovery interval value can be set here. An actively monitored network will need the IP discovery polling interval to be set to at least once a week, but because of time restrictions of a normal NNM system audit, the discovery interval will need to be set much lower. The NNM administrator should be with the auditor as they decrement the discovery value by one day at a time, once an hour during peak network usage time, until the discovery interval is within the time the audit's scope has allowed for the system audit. The IP Discovery interval should not go below once a day for a Class A network. The command *ovstatus –v netmon* should be run again to see what the accelerated device discovery has done to NNM's ability to keep up with the other polling activities it must carry on accurately. The auditor needs to pay careful attention to whether the system

9

is getting behind in its polling and DNS lookups. New nodes that are found will pop up in the maps, to be either: managed, unmanaged, or hidden after their SNMP configuration is updated and secured.

Objective/Subjective: The results of this test are objective. If new SNMP-enabled nodes are placed on the monitored network, they will appear on the NNM maps. The NNM administrator will then need to meet with that device's system administrator and have the SNMP settings reconfigured to conform to the corporate SNMP policy.

Step 2: Review SNMP RO string corporate policy

Reference: http://www.sans.org/rr/papers/44/376.pdf

Control Objective: This is necessary to ensure that the network is configured securely for SNMP-enabled devices. The need for an SNMP policy is often overlooked, until there is a successful attack on the network after system information was pulled off of an SNMP-enabled device.

Risk: If a device was put out on the network that was configured to use an RO string for monitoring purposes, but the string was easily guessed or brute-force cracked, the device could give system information to an attacker. This information could range from device specifics to network specifics depending upon the administrator or manufacturer default configuration. Tools such as SolarWinds' SNMP Brute Force Attack, can run through a dictionary of words and perform permutations on those words to guess the RO or RW string of an SNMP-enabled device automatically (http://www.solarwinds.net/Tools/Security/SNMP_Brute_Force/index.htm).

Compliance: The RO string that is chosen should be strong enough that it wouldn't be subject to a brute-force attack or easily guessed by the attacker. Strength can be difficult to test, but should conform to the enterprise's basic password policy.

Testing: The auditor should request the NNM administrator to produce the enterprise's password security policy that is enforced across the network. The RO strings that are used in the SNMP configuration window of the NNM RW map should pass the acceptable password security definition in the policy. If the RO strings do not pass, the NNM administrator should formulate an RO string migration plan for each group of devices, that are being monitored, that do not use an acceptable RO string.

Objective/Subjective: This test will provide subjective results, because it can be difficult to decide whether the RO strings that are being used on the monitored devices are strong enough. The enterprise's password policy should be used as a guide if there is no formal policy on SNMP configurations for network devices. This will ensure that reasonably strong RO strings are being used and enforced by the NNM system and administrator.

Step 3: *Public* should be removed as the SNMP RO string from all devices

Reference: http://xforce.iss.net/xforce/xfdb/6296

Control Objective: All SNMP-enabled devices usually ship from the manufacturer with *Public* as the default RO string, which should be changed or removed. This will verify that in the attempts to monitor mission-critical devices on the network using an NNM system, the NNM administrator has not opened up the possibility of an enterprise wide attack or successful reconnaissance mission using default RO strings.

Risk: If the default RO string is left on a device, an attacker can connect to that device using the RO string and view valuable configuration data. Each SNMP-enabled device on a network is configured to grant system administrators the ability to pull certain pieces of information off of them. If an attacker was able to connect to critical devices or networking equipment, intimate knowledge of the topology and configuration of the network would be known.

Compliance: The default RO string must always be changed in some way; either by removing the string completely, or changing the RO string to something other than *Public*.

Testing: On the NNM server, the RW map can be opened and the auditor can then go to Options > SNMP Configuration in the File menu. Under the Global Default tab, there is a text field that contains the RO community string; this should be something other than *Public*. The Set Community text field below should be empty for an NNM monitored network. Under the IP Wildcards tab, if there are networks explicitly defined to use a Community string other than the default global RO string, the devices on this network should be changed to use something other than *Public*. Under the Specific Nodes tab, if there are nodes explicitly defined to use a Community string other than the default global RO string, the device should be changed to use something other than *Public*.

Objective/Subjective: The results of this test are objective; the tests are straightforward and can be repeated on any NNM system being used for network device monitoring.

Step 4: SNMP RW string should be removed from all SNMP-enabled devices

Reference: http://www.sans.org/rr/papers/44/376.pdf

Control Objective: To mitigate the risk of having a device taken over remotely by an attacker, the RW string should be taken off of all SNMP-enabled devices on a network. The default RW string of *Private* should be removed from any SNMP-enabled device, but in an SNMP monitoring environment, there should never be RW strings on devices whether default or unique.

Risk: If *Private* is being used as the RW string on a device this needs to be removed, because an attacker could take control of the device with minimal effort. All SNMP community strings can be guessed using brute-force SNMP tools, which would find the RW string that is being used on a device. Tools such as SolarWinds' SNMP Brute Force Attack, can run through a dictionary of words and perform permutations on those words to guess the RO or RW string of an SNMP-enabled device automatically (http://www.solarwinds.net/Tools/Security/SNMP_Brute_Force/index.htm). Once an attacker knew the RW string of a device, they could not only take control of the device, but also learn valuable information about it and other devices that may have a trust relationship on the network with that device. Since this environment scenario is purely a monitored SNMP network, all devices should have had their RW string stripped off before they were added to the network.

Compliance: The auditor should consult the written SNMP policy, or discuss the verbal policy that is enforced. Stripping all RW strings off of all devices on the network should be one of the top priorities of the enterprise's SNMP policy. Reviewing the SNMP configuration of the NNM system may also reveal deviations

from an SNMP monitored network best-practice configuration.

Testing: The auditor should open the RW map and navigate to Options in the File menu and then select SNMP Configuration. This will open the SNMP Configuration window, where under each tab the auditor can verify that there aren't any RW strings defined globally, by subnet, or by specific devices. The auditor should also review the enterprise's verbal or written SNMP policy, and verify that it explicitly defines that all SNMP-enabled devices should be stripped of their RW string before they are placed on the network.

Objective/Subjective: This test will produce predictable results of either finding devices using an RW string or not finding any devices. What the auditor will find after reviewing the SNMP policy with the NNM administrator will also produce predictable results, because the policy will either cover RW strings being removed from devices before being placed on the network or it won't be covered by the current policy.

Step 5: Verify separate SNMP RO string designated for each group of devices

Reference: Personal Experience

Control Objective: This will add further security to the SNMP configuration of the network, which NNM needs to function safely as a mission-critical monitoring system.

Risk: If an attacker knows the SNMP RO string of a device, a simple SNMP utility such as SNMPwalk could be used to connect to a device, giving the attacker a lot of information about its configuration and current status. If each group of functionally related devices has its own SNMP RO string, then only the NNM administrator and the device's system administrator should know the RO string that each group uses. The group designation should be specified using the organizational groups within the enterprise itself. All of the servers that are being monitored should have an RO string that has been agreed on by the server team personnel, and the strength verified by the NNM administrator. All of the networking equipment that is being monitored should have a different RO string agreed upon within the group and then verified with the NNM administrator before changing the SNMP settings on all of the devices. The group division can be specified by the NNM administrator or Chief Information Office (CIO) of the enterprise, and each group's RO string should be kept a secret within the personnel group.

Compliance: There is a lot of lenience within the parameters of this test, but the auditor needs to check that there is a different RO string being used for each group of devices that are being monitored. A written security policy that has been approved by the CIO and NNM administrator should specify how often each RO should be changed on monitored devices and that each group should protect the secrecy of their RO string.

Testing: The auditor should discuss the written or unwritten policy concerning SNMP RO strings, and how well groups are adhering to the policy. An NNM map can be pulled up by the auditor, and then click Options > SNMP Configuration > Specific Nodes. This will give a list of devices that are being monitored with different RO strings from the default RO string. The *$NNMserver\conf\netmon.cmstr* file can also be pulled up within the NNM installation path on the NNM server. This will show the list of RO strings that NNM allows devices to be polled with and that can be used to discover other new SNMP enabled devices that may need to be monitored.

12

Objective/Subjective: The test results for this auditing step are subjective. There is a need for the outlined configuration, but may not be documented in the corporate SNMP policy. The appropriate procedure to accomplish this objective should be discussed and documented by the NNM administrator and CIO.

Step 6: Check devices from each group for SNMP configuration password

Reference: Personal Experience

Control Objective: This will verify that each personnel team that is in charge of a group of monitored devices has taken every security precaution to guard against their unique RO string being discovered or changed on devices.

Risk: Most SNMP-enabled devices can have their SNMP configurations changed remotely for ease of administration, but they are not always password protected by default. If the team that is in charge of a particular device updates the SNMP settings on the device to reflect having their team's unique RO string and remove the RW string, but forget to enable remote administration password protection, this device and others using the same RO string are at risk. If one of these devices is found by an attacker, potentially all SNMP configurations could be changed and active monitoring of those devices could cease to exist. The attacker would then also be able to get a wealth of information about the network itself, and the associated devices that use that RO string could then be compromised.

Compliance: There is not an exact science for this kind of random device testing, but if the auditor is able to find one device that does not prompt for a password before allowing access to the SNMP configuration of the device, then it was worth the effort. If no devices are found to grant this kind of access, then this test was still worth the effort to make sure that one device didn't disclose the RO string of many other similar devices.

Testing: If the auditor wanted to start by testing some random devices that are controlled by a network team within the enterprise; they can start by locating bridge-type connector symbols in the NNM RW map. This can be accomplished by selecting Edit in the File menu, selecting Find, select Object By Symbol Type, and then the Find By Type window will appear. A Connector can be selected under the Symbol Classes area, Bridge can be selected from the Symbol Subclasses area, the Apply button at the bottom should be pressed, and then select a random device from the list that NNM returns. Open a command prompt and telnet to that device, to test for a password authentication prompt. If a device that is not password protected can be found anywhere on the map, and the SNMP configuration can be accessed remotely using this method, the NNM administrator should notify that device's system administrator immediately to have all devices that carry the same configuration changed. The enterprise's SNMP policy should also be revisited to verify that it discusses the need to have any devices with remote SNMP configuration changing capabilities be password protected. Each group of devices should be tested, from networking equipment, to servers, down to print servers, and network printers.

Objective/Subjective: This test can be difficult to validate that all devices are compliant, but a brief check of some of the devices from each personnel team's list of administered machines can pay off. The amount of devices that are checked is subjective to the auditor's discretion of how compliant the NNM system has been

throughout other steps of the audit. This will dictate how many devices will be tested across the network of monitored devices.

<u>**2.2** - NNM System Remote Management Configurations</u>

Step 7: Verify that the NNM RW map has been secured

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This step will show that only the NNM administrator or an authorized NNM user can access and make changes to the RW map.

Risk: This test is very important due to the level of manipulation gained by having access to the RW map. Ideally the NNM server would be left logged in using a service account with a password that only NNM authorized users have. The screen would need to be left in a locked state in the server cabinet using the server cabinet's Keyboard Video Mouse (KVM) system. This will keep unauthorized users from logging onto the NNM server locally and running an *ovstop* command or *ovstop ovuispmd*, and then taking over the NNM RW map. Also, if an NNM RW map is not locked down in some form, anyone with access to a workstation that has an NNM remote client installation can change the configurations of the NNM system and learn the monitored networks' topology.

Compliance: This is not a binary compliance check; there are several ways to secure an NNM RW map, and the NNM administrator will have hopefully chosen a preferred method described below.

Testing: There are three ways to lock down an NNM RW map. The administrator can leave the RW map open on a locked screen in the server cabinet's KVM configuration. The RW map can remain open on the administrator's workstation at all times. The NNM installation path can be set to use NTFS authentication mechanisms, which would only allow RO privileges to users who aren't supposed to be able to change anything on the NNM RW map when using their NNM remote management client.

Objective/Subjective: This is a very subjective item in that there are choices as to how an administrator can fulfill this requirement, but one of these choices or a combination of them must be chosen.

Step 8: Check that NNM root directory Share specifies authorized NNM users

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This configuration will secure the NNM database, maps, and the overall NNM system from becoming corrupted.

Risk: If the database were to fall into the wrong hands, an inside attacker would know the entire topology of the enterprise, and could also manipulate certain events that they may have triggered in other malicious internal behavior. The NNM database could also become corrupted and the NNM system would go down. If the NNM maps were able to be seen or manipulated by an attacker, it could give away valuable network topology information, and also open the system to a possible misconfiguration. The configuration of the NNM installation could also be damaged by allowing access to the installation path on the NNM server.

14

Compliance: This check will not necessarily provide a yes or no answer, because the administrator would need to be very strict in what access is given out on the NNM system. If the NNM administrator trusts certain team members to have a specific level of access to the installation path, the administrator needs to be aware of the amount of access that configuration will bring to a team member.

Testing: The NNM installation path would need to be checked for compliance to the item described. Right-clicking, selecting Sharing, and then the Permissions button on the root installation folder of NNM, will show which users have been setup with file access. The NNM administrator would need to make a list of users and their privilege level in this folder, and then the auditor can check for compliance. The administrator will need to be reminded that this configuration not only effects whether or not a user will have access to the NNM .conf files, database, and topology files, but also that the configuration decides what level of control the NNM remote management clients have over the NNM maps.

Objective/Subjective: This test will produce objective results for the auditor. The list of authorized NNM users and their access-level can be compared to the list of users allowed to the root directory Share.

Step 9: Lock down web access to specific authorized NNM users

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This will check to make sure that only the authorized NNM users that the NNM administrator wishes to allow access to view the NNM web maps have that ability.

Risk: It is very important to make sure that the NNM web maps are not accessible to all users within the enterprise. Only authorized NNM users and possibly management would need to see the web maps. If an inside attacker was able to gain access to the topology of the LAN/WAN, they would have an exact map of mission-critical devices that could be attacked.

Compliance: The auditor will check the NNM web map security precautions that have been taken by the NNM administrator. If certain NNM configuration files have been created and setup correctly, then the NNM administrator will have restricted NNM web maps to only be accessed by authorized NNM users..

Testing: Look for *$NNMserver\www\etc\htpasswd* to see if the NNM administrator has modified the User Authentication Password File with authorized NNM user names. The auditor should then point a web browser to http://$NNMserver/OvCgi/ovlaunch.exe and see if they can connect. The auditor should not be able to connect without a valid username and password.

Objective / Subjective: This test is fairly objective in that most of the time an auditor will be able to tell fairly quickly as to whether or not the NNM administrator has locked down his list of authorized NNM users and their privileges concerning web map access.

Step 10: Web map access should be logged in NNM *audit log*

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This step will verify that all NNM authenticated user's web map access is being logged in the audit log within the NNM system.

15

Risk: This check addresses the need for the NNM administrator to know exactly who has been accessing the NNM web maps. The audit log will show which user, from which host, at what time, and which URL's were accessed during the login period. The list of URL's will specify exactly which programs were utilized from the NNM launcher applet window.

Compliance: This test can be done once the auditor gains access to the NNM installation path on the NNM server, to verify the logs are being modified correctly by the NNM system.

Testing: The Login Log should be opened at *$NNMserver\www\logs\login_log*, and the Access Log should be opened at *$NNMserver\www\logs\access_log*. The auditor should point a web browser to http://$NNMserver/OvCgi/ovlaunch.exe and see if they can connect with an invalid username; this will give an Access Denied entry in the *login_log* file. The auditor should try this again, and use a valid username; this should give an 'allowed entry' in the *login_log* file. This entry will also have information about the host, date, and session number as well. The *access_log* file will have an entry with the same session number, and will also itemize the URL's that were accessed during the session.

Objective / Subjective: This test will produce objective results that can be verified quickly. If the NNM web map log files are being appended with new events, then the NNM system is setup to log web map access.

Step 11: Log user and restrict access to certain web tools in launcher

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This test verifies that authorized NNM users are restricted to RO or RW map access. The authorized login of a user should be enforced for each session, and an authorized session should timeout at a specified interval.

Risk: The risk that is being mitigated with this configuration is the chance that an unauthorized user could open a web browser and navigate to the NNM launcher and open a web application. The session configuration file should be configured to force a login session for each application URL accessed from the NNM launcher. The NNM administrator should also enable login logging and access logging. The NNM administrator should use this configuration file to set a session timeout value for the allowed number of hours before re-authentication of an authorized NNM user.

Compliance: There is a lot of functionality to check within this file, but it doesn't take long to check the compliance of the file. If the NNM administrator wishes to have maximum security for the system, this file should have all options set to 'on', and a timeout value set lower than the default value.

Testing: Look for *$NNMserver\www\conf\session.conf* and see if the file has been modified by setting all login fields to 'on'. The session timeout should also be set to a timeout value below the default of 9 hours. If it were set at four hours, an authorized NNM user would need to login in the morning, and it would ask for another login in the early afternoon. This would give the NNM administrator two traceable logins for each authorized user each day.

Objective / Subjective: The test results for this auditing step can be somewhat subjective, because the session timeout value may need to be set at a different interval for different scenarios. Since there is no best-practice documented for NNM

16

monitoring implementations, a timeout value of four hours will be used for testing.

Step 12: Define web access group restrictions for authorized NNM users

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: Check to make sure that authorized NNM users have been placed into web access groups that assume predefined access-level privileges into the NNM system.

Risk: The auditor will want to make sure that only certain authorized NNM users will be able to use certain tools from the NNM launcher. There are seven group-level privileges that can be assigned to users. Certain levels grant access to NNM web tools that only certain authorized NNM users will need. Some tools will give an SNMP report of an entire device or of the NNM system itself, which would give a list of the current and past status polls of the entire monitored network.

Compliance: The auditor should check to make sure the authorization file is currently in use, but the group level privileges/restrictions are not documented by HP as to which tools can be used by each group. This check will take some time to decide whether the configuration is correct or not. The NNM administrator may also have to elaborate on the list of authorized NNM users as to their allowed web tool access from the NNM launcher.

Testing: Look for *$NNMserver\www\etc\htgroup* and see if the file has been modified by putting authorized NNM users into groups with more permissions than specified by the NNM administrator. The file syntax should also be checked to make sure that there are no + signs next to group names. The auditor will need to be added to the list of authorized users in the authentication file that was discussed earlier. The auditor's login credentials should then be placed into one of the seven authorization-file groups. The auditor can then log in and document which tools they can access. This test will be repeated until the auditor has been placed into each group, logged in, and then the available tools have been documented. The auditor will then need to go through the list of authorized users to decide whether the group that each user has been placed into is actually the ideal group for that particular NNM user.

Objective/Subjective: This test will be fairly subjective due to the amount of possibilities for each user's allowable access. These are always replicable tests, but will definitely be different for each enterprise.

Step 13: Verify that NNM users aren't using NT domain passwords

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This check will make sure that authorized NNM users set their web access passwords as something other than their NT domain passwords, since web map authentication sessions send passwords as cleartext.

Risk: If an attacker placed a packet-sniffer such as WildPackets' EtherPeek on the subnet where the NNM server resides, the NNM web access password of an authorized NNM user could be intercepted during an authentication attempt to the NNM system. Once the attacker discovers an authorized username and password, a web browser can be used to navigate to the NNM web server and log in as an authorized NNM user. This will look like an acceptable web session in the NNM logs, and may go undetected for some time. The attacker would then know the topology of

17

the monitored network, and could use as many of the NNM launcher tools that were allowed to that authorized NNM user against SNMP-enabled devices. The packet sniffing and password capture is virtually unavoidable to accomplish this extreme test, but the risk is that some authorized NNM users may choose to use their current NT domain password. If this password were intercepted, the attacker would then have the capabilities of logging into any device on the network using an authorized domain user.

Compliance: An auditor might feel this test is too difficult and isn't worth the time it takes to accomplish it. Although, an enterprise may be concerned enough with the passwords that are used by authorized NNM users to go through with this test. When the NNM administrator originally sets up the new authorized NNM web user in the *htpasswd* file on the NNM server, each user should be reminded that they are not to use an NT domain password. Compliance must be tested secretly so that an authorized NNM user is not compelled to change their domain password right before the test.

Testing: The auditor would need to setup an Ethernet packet-sniffer on the subnet where the NNM server resides, and capture traffic in the morning when the users would be signing on. Successfully sifting through the traffic to intercept an NNM web access password may be too difficult on a busy network subnet, but the *ovuispmd* service could be administratively restarted on the NNM server, which would force the users to re-authenticate. This will help the auditor be able to find the password, because they will know a timeframe to help pinpoint the password in all of the traffic. The intercepted credentials can then be compared to the current NT domain password that is being used by the authorized user, by gaining access to the NT User Manager.

Objective/Subjective: This step is an objective test, but very time consuming and the NNM administrator will need to decide if it is worth the time to verify compliance of the authorized NNM users.

2.3 - NNM System Software Configurations

Step 14: Verify that management stations and consoles match NNM versions
Reference: Personal Experience
Control Objective: Check that all authorized NNM users are using the same NNM remote management client version as the management station has installed on it.
Risk: If the NNM remote management clients that are installed on authorized management consoles access the management station with a client that does not match the version of the NNM software that resides on the management station, the NNM database could become corrupted. This may not be a problem with patched releases within major versions, but this does not comply with best-practice. The NNM administrator should not allow authorized NNM users to connect to the management station with versions of the remote client that differ from major version releases. A filesystem Share can be setup on the NNM server that houses the latest release of the remote management client for authorized NNM users to download and install. This NTFS Share should be locked down to RO privileges and to authorized NNM

18

users only.

Compliance: The auditor can view the version numbers on the management station and each management console, and compare version numbers to check for compliance.

Testing: The auditor can open the map on each management console and compare the version to the management station. In the File menu, select Help, and then About HP OpenView. The second line on this window will state the version of the client at the end of the line that states: HP OpenView Windows NNM Release. The version string (0X.XY) on the management console should match the X's on the management station, but may differ on the Y if patches have been applied to the major release installation of NNM.

Objective/Subjective: The results from this test are objective and easy to verify. The version number must be located on the management station and each management console must match it.

Step 15: Verify the NNM management station is running the latest NNM release

Reference: Personal Experience

Control Objective: This will check to make sure that the NNM system is running the latest release of the NNM software.

Risk: If the NNM system is not running the most current release of the NNM software, then the system could be at risk of falling victim to unstable code in the older software. This could cause a system outage, loss of device event data collection, or a loss of network administrator device status notifications. The older NNM software could also be susceptible to an attack due to code vulnerabilities being exploited by an attacker. On the other hand, the NNM administrator may not want to use the latest release of the NNM software due to possible bugs in the new code, or because of new recommended hardware specifications the company's budget may not be able to meet.

Compliance: This is an easy test for the auditor, but there are other enterprise specific considerations that must be made by the NNM administrator after this step has been completed. If the NNM system is not running at the most recent release reported by the NNM website, the NNM system will fail this test.

Testing: The auditor can simply point a browser at http://openview.hp.com/products/nnm/index.html to check for any new major releases that are available. A new release upgrade can only be accomplished if the enterprise has a software subscription license for NNM with HP. The management station map can be pulled up to compare the current release version to what is running on the management station. Once the map is pulled up the auditor can navigate to the Help button in the File menu and then choose About HP OpenView from the drop down list. The second line on this window will state the current version of the NNM system at the end of the line: HP OpenView Windows NNM Release. The X's in the string (0X.XY) compose the major version release number.

Objective/Subjective: The testing results produced will be objective. This is an easy test for compliance, but other enterprise specific constraints may complicate the decision for the NNM administrator to upgrade the NNM system or not.

19

Step 16: Verify NNM backups run locally, weekly, and are stored remotely

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This is to verify that in case of a hardware failure or a physical disaster where the NNM server is housed, that the NNM system can be restored to its last known good-state from backup.

Risk: If there were no backups made of the NNM system, or the backups weren't taken to a remote site that would be safe during a disaster, the NNM administrator would have to work frantically to build a new Windows server, reinstall NNM, and reconfigure all of the settings. This would disallow all active SNMP monitoring of the mission-critical devices on the network. Valuable forensic data from the NNM event database would also be lost and unrecoverable in both instances. If NNM is used in a NOC scenario where SLAs depend on constant device monitoring to ensure network uptime, there could be a huge financial penalty weighed against the enterprise for lower than average uptime of the network devices that are being monitored.

Compliance: This can be verified easily, by stepping through the backup procedures outlined in the testing are below. This is a very simple and repeatable test for a very complex and valuable system.

Testing: A discussion should first take place between the auditor and the enterprise's backup solution system administrator. The enterprise's written backup policy should be reviewed to ensure that the NNM system is being backed up at least once a week. The backup policy should also define a reasonable disaster recovery plan. The backed up NNM system data should be offsite in a different physical location of the enterprise's campus. The time of the NNM system backup can then be determined by checking the IP or DNS name that is set in the backup system's configuration. The NNM server data backup should happen directly after the *ovbackup.ovpl* script has run within the NNM system. The auditor should log into the NNM server and verify that there is more free space than used space left on the hard drive, which will ensure NNM server stability after backups have been run locally of the critical NNM system configurations. The Scheduled Tasks window should then be pulled up by going to Start > Settings > Control Panel > Schedule Tasks. There should be a scheduled task that runs the *ovbackup.ovpl* script at least once a week, and directly before the enterprise's backup solution does the weekly backup of the NNM server's hard drive.

Objective/Subjective: This is an objective test that is common for an enterprise backup policy concerning mission-critical servers. The tests will product predictable results, and are very important for aggressively monitored networks that must always be up.

Step 17: Check NNM patch level for compliance with most recent release

Reference: HP - OpenView Managing Your Network for NNM

Control Objective: This step will ensure that the NNM system has the latest patch release installed, which will correct discovered vulnerabilities or bugs in the NNM system code.

Risk: If a vulnerability is discovered within NNM, an attacker may be able to exploit a hole within the NNM system code and compromise the integrity of the NNM database. The attacker could also discover information about the SNMP configuration of devices being monitored by the NNM system. Some NNM patches are released to correct a

bug in the NNM system code, which is either a security concern or the system needs added stability in the NNM software. In some enterprise's if a link or device on the monitored network were to go down, because the NNM system stopped functioning and no one was notified, SLAs with clients could be broken and financial penalties would be assessed.

Compliance: The NNM management station will house the information that shows the current patch level of the NNM system. This can be verified by navigating to files that are installed during a patches installation and reviewing the version.

Testing: The auditor can log onto the NNM server and open the RW map. While the map is being loaded, the splash screen will show the current consolidation patch level of the system. This screen can also be brought up by navigating to Help in the File menu and then selecting About HP OpenView. The second line of text will state HP OpenView Windows NNM Release B.0X.XY. The X's are the major release version number, and the Y value is the current consolidated patch level of the system. There are other patches that are released that will not affect the Y value when applied to the system. Most of these patches are quick fixes for a discovered bug or vulnerability, and after enough of the intermediate patches are released there is a consolidated patch that rolls up several new fixes, but also holds all of the old ones as well. These affect the value that is represented by the Y, and all subsequent patches that are released will be functionally dependent upon the prior consolidated patch and can't be installed without them. The site

http://support.openview.hp.com/cpe/patches/nnm/6.4x/win.jsp should be referenced to find out the latest consolidated patch release number. To test whether or not the latest consolidated patch has been applied to the NNM system, the auditor can navigate to $NNMserver\Patches\Patch_Name\patch.txt. If this file does not exist, then the patch has not been installed. The site

http://support.openview.hp.com/cpe/patches/nnm/6.4x/win.jsp should be referenced to find out the latest intermediate patch release number. To test whether or not the latest intermediate patches have been applied to the NNM system, the auditor can navigate to $NNMserver\Patches\Patch_Name\patch.txt. If this file does not exist, then the patch has not been installed.

Objective/Subjective: This is an objective test to see whether the NNM system is at the current patch level. The test will produce repeatable results that will compare the NNM system's current patch version to the most recent patch number on HP's NNM support page.

Step 18: Review corporate policy used for NNM system patching

Reference: Personal Experience

Control Objective: This test is to verify that a corporate policy describing the procedures to patch mission-critical servers is being used by the NNM administrator concerning NNM system patching.

Risk: If an NNM system vulnerability or bug is found in the code, then an intermediate patch will be released by HP. These patches should be installed as soon as possible to ensure maximum uptime for the NNM system. The intermediate patches are not as dangerous for the NNM administrator, because the patch does not usually call for the NNM server to be rebooted after installation. Certain precautions

21

should be addressed by the NNM administrator as to the patching procedure that will be used for the NNM system. These precautions must be taken into consideration to prevent NNM system file corruption. Even though NNM system patches have an uninstall feature on them, NNM database corruption is possible and some files that may have been customized by the NNM administrator could be overwritten. The NNM system processes will also have to be stopped temporarily until the installation has finished for both intermediate and consolidation patches.

Compliance: The NNM administrator either will or will not be able to produce a written corporate policy concerning mission-critical servers that is being used for guidance with NNM system patching.

Testing: The policy for patching the NNM system should be inline with the corporate policy of patching mission-critical servers. The document should state when the NNM system is allowed to be taken down for maintenance, who is allowed to apply approved patches, and what steps must be taken before a patch can be installed. For some enterprise's the network monitoring server is considered mission-critical, so network monitoring outages should be approved by a person in management higher than the NNM administrator. Also, a list of items to be checked before and after the patch is applied should be documented. If a patch has been released for the version and platform of NNM that is currently running in the enterprise, the NNM administrator should only apply the patch after the latest NNM system backup has been run successfully, and it has been physically backed up. Also, the patch should not be applied during a network utilization peak. Any consolidated patches that other intermediate patches are dependent on must be verified to have been installed before the new intermediate patch is applied. All RW and RO maps must be closed on the management station and consoles, which can be done by running *ovstop* at the command line on the NNM server. The SNMP Master Agent and Adapter for NT services need to be stopped, by going to Start > Programs > Administrative Tools > Services, right-clicking each service, and selecting stop from the menu. During patch installation, the setup program will prompt the administrator to either save the current status of the NNM system or overwrite and add files without saving the current settings; the administrator should always choose to save the original settings, so that the patch can be uninstalled later if the database becomes corrupt or the NNM system is not functioning properly. Once the installation prompts for a reboot, the administrator should allow it to reboot. When the NNM server comes back up, a command line should be brought up and the command *ovstatus –v* should be run. This will list all of the services that are needed by the NNM system, and each service should show a status of Running. If one of the services does not start back up, the command *ovstart* service (name of service not running) should be run; this will ensure all necessary services are running and functioning correctly. The SNMP Emanate services should also be checked to see if they are running. The NNM administrator can then navigate to Start > Programs > HP OpenView Patches > HP OpenView NNM 6.4 and verify that the patch name and number shows up in the list. One final test is when the RW map is pulled up after reboot, the NNM Administrator can verify that the version number changed to B.06.41 if the patch that was installed was the first consolidation patch for that particular NNM system major release.

Objective/Subjective: This test will produce subjective results. The auditor will have

to make a judgment call as to whether there is compliance with this step. Consideration will have to be given to the fact that the NNM administrator may be handed down corporate policy that specifies when and if the latest patches can be placed on the mission-critical equipment.

Step 19: Check for optimal hardware configurations on the NNM server

Reference: Personal Experience

Control Objective: This test will check for any hardware misconfigurations or NNM server settings that could be changed to help make the status of the monitored devices more accurate, and help alleviate as much stress as possible from the NNM server's hardware.

Risk: The hardware configuration of the NNM server must comply with the testing steps below, or the status of certain mission-critical devices may not be portrayed as accurately as a Network Operations Center (NOC) may demand of the NNM system. The more stress the NNM software is putting on the NNM server, the easier it would be for an attacker to successfully perform a Denial of Service (DoS) attack on an already heavily taxed NNM server CPU.

Compliance: The quick checks that are outlined below can prove to be very helpful in solving possible performance issues that may arise on a heavily utilized NNM server. The server will either conform to these best-practice tests or it will fail.

Testing: An NNM server's hardware should be at least double the minimum hardware requirements that are defined in the ReadMe file distributed on HP's install medium. This file can be found in the root directory of the installation CD for the NNM system. The first page in the HTML document will specify what the minimum system requirements should be for the version of NNM running on the particular platform being used. NNM 6.4 specifies 512MB of RAM, 1GB of free hard drive space on an NTFS partition, and an Intel Pentium 333Mhz processor. The hard drive should be much larger than double the 1GB minimum, since the NNM system backups will reside on the local hard drive after running each week. The auditor should right-click on My Computer and select Properties. This window will describe the current hardware configuration on the NNM server under the heading Computer. This window can be closed and then double-click the My Computer icon. The subsequent window will hold the icons for the local disks, and mapped drives. The local disk the NNM installation resides on, should be selected and checked for free space conformance. For best-practice there should be at least double the amount of free space on the drive as there is used space, due to the size of NNM system backups that will need to be performed weekly. The Network Interface Card (NIC) used by the NNM server to connect to the network, should have a 100 Mbps, Full Duplex network connection. Someone from the enterprise's network operations team will have to verify these settings on the network switch that is directly connected to the NNM server NIC.

Objective/Subjective: This test will produce subjective results. Not all environments will need the hardware specifications above, but most aggressively monitored networks will need an NNM server with that amount of hardware power.

Step 20: Check for optimal software configurations within the NNM system

23

Reference: http://openview.hp.com/sso/ecare/keyword

Control Objective: There are some very important configuration settings that are only discussed on the OpenView technical knowledge base on HP's site, which some less experienced NNM administrator's need to be aware of before an NNM system deployment.

Risk: If certain settings aren't configured on the NNM system properly, the status of certain mission-critical devices may not be portrayed as accurately by the NNM maps as a NOC may demand of their NNM system. A poorly configured NNM system that is heavily using the NNM server CPU will also be more at risk in falling victim to a DoS attack launched at the NNM server.

Compliance: This test can produce results that are difficult to interpret and may dictate a judgment call when deciding on a pass or fail decision for the test. The outcome of this test depends greatly upon the NNM administrator's personal preferences when configuring the NNM system, the corporate network monitoring policy that has been handed down to the administrator, and the hardware constraints the administrator has to work with.

Testing: The Network Polling and SNMP Configuration areas on the NNM RW map contain the most important configuration areas of the NNM system, which can make or break the monitoring software's functionality. A CPU check on the NNM server must first be taken into account. This can be accomplished by logging onto the NNM server, hitting *CTL+ALT+Del*, and then selecting Task Manager. When the Task Manager window comes up, the Performance Tab should be selected and the current CPU load should be monitored for at least a minute to verify that the CPU never hits 100% load. The auditor should open the RW map and navigate to Options in the File menu and then select Network Polling Configuration. The configuration of these options depends greatly upon how aggressively the NNM administrator feels the enterprise should be monitoring their mission-critical devices on the network. The option to perform topology checks and configuration checks should be marked, and the time values should be set to at least once a week. The CPU load should be monitored during any changes, as the NNM server may become overloaded during peak network utilization, or because the NNM server's hardware has hit its processing capacity. The IP Discover tab should then be selected and the Discover New IP Nodes checkbox should be marked. If the Discovery Polling Interval is set to auto-adjust and the CPU load is still relatively small, then it is probably fine to leave it there; by default this is set to rediscover the entire network within 24 hours. This is an area though, where the NNM administrator may be able to take a little of the strain off of the hardware by setting the NNM system to a fixed discovery interval. The interval value should be set to at least once a week, unless the NNM administrator isn't worried about catching unauthorized devices that are being placed on the network with SNMP services misconfigured. Finally, under the Status Polling tab, the Perform Status Polling checkbox should be marked. From this area you can open the SNMP Configuration window by clicking the Configure button at the bottom of the window. The node status polling interval will probably not be able to be set lower than once every 5 minutes, and the CPU load should be monitored during any of these changes to ensure it never reaches 100% and stays for an extended period of time. The timeout and retry values for the node polling configuration should stay at the default

24

value, unless a device's status is not accurately portrayed by the NNM maps. This could also be attributed to the link to that device being saturated and the polling traffic is being dropped for higher priority traffic on the line. The NNM system could also be running behind with its device status polls, because it can't poll all of the devices as quickly as the interval value specifies. If this happens for a particular subnet, the subnet should be added as an entry into the IP Wildcards section of this window, by using a star to contain any devices within a certain subnet that should by treated with a different timeout and retry value. If there is a problem with a particular device only, this same configuration can be done in the Specific Nodes area of this window.

Objective/Subjective: These tests will produce somewhat subjective results. The testing steps must be followed carefully, but there are many possible variables in this scenario that should be taken into consideration by the auditor when submitting the test results.

Step 21: Check for commonly overlooked NNM performance issues

Reference: Personal Experience

Control Objective: This is to verify that the NNM system is functioning optimally with the given hardware and software configuration. Network monitoring software must be running at peek performance levels to be effective in alerting a device's system administrator to check on the status of that device.

Risk: There is a risk that an NNM system may go unchecked for weeks at a time. Because if the NNM system is functioning it doesn't mean that it is functioning as well as it could be, or may not reach an unstable or unreliable state in the near future. A poorly configured NNM system is more at risk to a successful DoS attack, because of its already fragile operating state.

Compliance: This is a yes or no list of checks, which must be done to verify NNM system optimal performance. The testing steps outlined below discuss exactly what the auditor should see after the test commences, anything that deviates from the predicted response is a failed test.

Testing: The first test will be for CPU load. This can be accomplished by logging onto the NNM server, hitting *CTL+ALT+Del*, and then selecting Task Manager. When the Task Manager window comes up, the Performance Tab should be selected and the current CPU load should be monitored for at least a minute to verify that the CPU never reaches a 100% load. The performance of the NNM system's primary purpose can be measured by a graphing tool built into the maps. The auditor can open the NNM RW map, select Performance from the File menu, and then select Network Polling Statistics. The subsequent window will have a real-time graph that shows the status polls and SNMP polls that are running in the background to the monitored devices. All of these graphs should stay above zero the majority of the time; otherwise, other network polling configurations may need to be looked at to be reconfigured. Both the Task Manager and the Network Polling Statistics graph should be left open and watched by the auditor to verify that the graphs stay in the acceptable ranges decided upon in the previously discussed steps. The auditor should then open a command prompt window on the NNM server, and type in the command *ovstatus –v*. This will list the NNM processes that should be running. The list should be scanned to verify all processes are up and running. If one of the

25

processes is not running for some reason, the command *ovstart –c service* (the name of the service not running). The command *ovstatus –v netmon* should then be ran in the command line. The output we are concerned with here is whether or not there are overdue polls and how far behind they are. As before, if the NNM system is falling behind, certain Network Polling parameters may need to be reconfigured to allow the network to be adequately monitored. Another concern is whether or not the NNM system is having problems doing name resolution for devices on the maps. If this command gives us output that details multiple Domain Name Services (DNS) requests with high response time, the NNM administrator may need to setup a local DNS server on the NNM server. The NIC would then use the local DNS server as the primary and the enterprise's primary DNS server as the NNM server's secondary DNS server. A final step that is sometimes overlooked is that of the speed of each device when answering the NNM system during an SNMP configuration check. If a device makes the NNM system wait to complete a scheduled poll, it will be during a full SNMP configuration check. If a device stalls during an SNMP configuration check it is usually during the routing table call. Routing tables are called to aid in IP device discovery and topology configuration changes that should be reflected on the NNM maps. This will only be a problem on routers with large routing tables, which are possibly connected to the network with slow links, and little CPU to spare. Border routers are usually the culprit, and can take many minutes and sometimes hours to dump out their routing tables to the NNM system. Because border routers are at the edge of a network these routing tables aren't always that important for IP device discovery on the internal network. These devices should be selected by the auditor on the RW map, Fault should be selected on the File menu, Network Connectivity, and then Poll Node. This will open the SNMP polling window which will run through an ordinary poll step by step, and the system will probably stall when trying to retrieve the routing table. This can cause unneeded stress on the border routers, and cause the NNM system to become behind in its polling device list. The NNM administrator may want to consider specifying within the SNMP Configuration area that this device does not need to be a part of a full SNMP configuration check.

Objective/Subjective: These steps produce objective results and must be done to ensure the NNM system is running at peek performance levels. The commands discussed will produce verifiable NNM system output.

2.4 - Physical Security and Login Access

Step 22: The NNM server should be in a locked cabinet with backup power

Reference: http://www.activsupport.com/network/vpn_security/physical_security.html

Control objective: Verify the physical security and backup power plan for the mission-critical NNM server. In an enterprise's communications room there will usually be many rows of locked server cabinets drawing power from redundant power circuits. These power circuits should have a backup power supply in case of a power failure to the building housing the NNM server.

Risk: This check will show that only authorized personnel can gain access to the NNM server and that it has adequate backup power in case of a failure. This will

prove that there is not a physical security concern for the system. A physical breach of security can be fairly prevalent in an enterprise that is not concerned with malicious inside attackers.

Compliance: This test results for the NNM server's physical security are binary, and will prove whether the NNM administrator has complied with the details of this audit step or not. The test for adequate backup power may be more difficult to reach a decision for the auditor, since it can be viewed as a judgment call.

Testing: A physical survey of the server location and power architecture will verify that the server resides within a limited access room, within a locked server cabinet, and with sufficient backup power. With the system that is being audited, there is also a concern about the physical security of the external modems used for the UMS. These should also be found and shown to be protected by a locked cabinet.

Objective/Subjective: The test for physical security is objectively verifiable. The backup power architecture compliance may be difficult to determine and subjective in nature. A campus operations employee may need to be consulted to verify testing compliance.

Step 23: Check management console in NOC for login credentials used

Reference: Personal Experience

Control Objective: This test will verify that a special service account has been created and been used for the NOC's NNM management console installation. If an NNM system is being used purely to monitor a network, depending on the size of the enterprise, there would probably be a NOC of some sort, which would have an NNM RO map up 24 hours a day for monitoring.

Risk: A domain account would need to be used to log into the machine where the NOC's RO map resides for active network monitoring. There is a security concern when using a domain account that is logged into a computer in a common area that never has the screen locked. If an employee wanted to do something inappropriate on the network without their credentials being tracked back to them, they could use the computer that always has the screen unlocked and is logged on with someone else's credentials. That is why a service account with limited capabilities should be created for the sole purpose of being used to log into the machine that holds the NNM RO map in the NOC. If this computer was setup any other way, illegal behavior could be performed on that machine and it would be traced back to the user who logged into it, and probably wasn't even around at the time. The service account username would also need to be allowed read-only access to the NNM installation root directory on the NNM server to be able to connect from an NNM remote management client. This account should not be made an Administrator on the NNM server, or someone could take control of the NNM server from the NNM RO map computer in the NOC.

Compliance: A configuration inspection on the NNM RO map computer in the NOC will produce repeatable and expected results for the auditor's testing. The auditor will check the NNM server Share configuration and the service account must only have read-only access to the filesystem Share or the NNM server will not pass the test.

Testing: The auditor should first look on the NNM RO map NOC computer to find the username of the user account that is currently logged into the computer. This name should then be taken to a person in charge of NT domain accounts and they can

27

determine the rights this user has on the network. This should be a service account with minimal capabilities on the network, with the NNM administrator as the owner of the account. This service account should not have Internet access, an email account, or remote network access capabilities. The auditor should then log onto the NNM server and verify that the service account is not an administrator on the NNM server. This can be done by right-clicking on My Computer, selecting Manage, expanding Local Users and Groups, selecting Groups, and double-clicking Administrators. The auditor then needs to navigate to the NNM installation path and verify that the service account is only allowed read access within the installation Share. Once navigated to the root installation directory, the auditor can right-click on the NNM root directory and select Properties. The Security tab should then be selected, the service account should be selected from the list of authorized users, and then verified that this account has read-only access.

Objective/Subjective: This test is repeatable and will provide objective results. The enterprise may or may not have a computer in the NOC dedicated to keeping an NNM RO map open for quick reference, but if so, these tests are important.

Step 24: Check to see who has administrative privileges on the NNM server

Reference: Personal Experience

Control Objective: This will take into account several concerns for the security of the NNM server, and the ability of an attacker to get around certain security precautions that HP has put into place for NNM.

Risk: This item is sometimes overlooked, because of the lack of time HP spends discussing it in the NNM manuals. The NNM administrator needs to make sure that he knows exactly who is an Administrator on the NNM server. If an attacker had administrative level privileges on the NNM server, then they could remotely connect to the NNM server using Microsoft Terminal Services or Computer Management, and stop and restart the services. An attacker could then use an NNM remote management client to view RO or RW maps. The *ovuispmd* service could also be stopped and restarted, and unless an NNM user was currently working on an NNM map this could go unnoticed for quite a while. This risk is seen only if there aren't any Permissions set on the root directory of the NNM installation that specify administrative RO or RW user privileges.

Compliance: This can be verified quickly using the list of users that the NNM administrator created outlining the authorized NNM users that are allowed access to the NNM maps and their specific level of privilege. The auditor can compare the list of users that should be setup as OS Administrators to the list on the NNM server.

Testing: The auditor can use the list of authorized NNM users and verify using Microsoft Computer Management that only NNM administrators are also Windows administrators on the NNM server. The auditor will then try to connect remotely using Microsoft Terminal Services or Computer Management and see if a user not specified as an NNM administrator can connect to the NNM server. The auditor will also need to determine whether an unauthorized NNM user account is able to connect to the NNM installation path using an NNM remote management client and is able to open an NNM map.

Objective/Subjective: This test will provide the auditor with independently verifiable

28

results. The list of authorized NNM users can be compared to the list of OS Administrators on the NNM server.

<u>**2.5**</u> - <u>NNM Server Operating System Configurations and Hardware</u>

Step 25: Login access to the NNM server should be logged locally by Windows

Reference: http://my.brandeis.edu/bboard/q-and-a-fetch-msg?msg_id=0000Ya)

Control Objective: This will check to make sure that the NNM server is keeping track of who has successfully and unsuccessfully logged onto the NNM server.

Risk: The NNM administrator needs to be able to track who logs into the NNM server. Once the NNM system is in a known good-state, any adverse changes to the NNM system need to be traceable back to a known domain user. The NNM administrator also needs to be aware of anyone trying to brute-force into the NNM server either remotely or a local login. The failed and accepted login attempts will be stored in the Security Tab of the Event Viewer, and a large amount of logins in a short amount of time will let the administrator know the NNM server is being attacked. If changes are made to the NNM system by an authorized NNM user with administrative privileges, the NNM administrator will know who made the changes and can speak with them about future maintenance on the NNM system and server.

Compliance: This step is to verify the NNM administrator has auditing capabilities setup on the NNM server to log successful and failed login attempts. This test has a binary result; the NNM administrator will either have the NNM server configured to audit logon attempts or not.

Testing: The auditor will logon to the NNM server, and then go to the Event Viewer under Start > Programs > Administrative Tools. The security tab in the Event Viewer will most likely be empty unless the NNM administrator has explicitly defined a security auditing policy elsewhere. If there are results in the security tab, then the event ID column header can be selected to sort the events by ID number. If there is a match between 528 and 547, then the NNM server is logging certain logon attempts.

Objective/Subjective: This test is repeatable and will produce objective results. The test steps through the configuration checks that must be used by the Windows log files.

Step 26: Event Viewer local log settings should be reviewed

Reference: http://support.microsoft.com/?kbid=320121

Control Objective: The NNM server logs should be set to not grow past a set size limit. The logs should also be configured so that they overwrite the oldest log entries, and aren't limited to a length of time that the entries are stored on the NNM server.

Risk: If the NNM server logs are left at default settings, the logs may not be allowed to grow large enough to still hold relevant data. If the Windows server that the NNM system resides on generates a large amount of system events that are stored in the logs, it may be necessary to increase the allowable maximum log size. If the maximum log size isn't set larger, then the oldest events will be trimmed off of the end of the file too quickly for forensic analysis if needed on the NNM server. There are three choices to overwrite (trim) the oldest events off of the end of the log file, but by

29

default the log file is set to trim off any event that is older than seven days once the maximum file size has been reached. This setting will usually trim off a large amount of events that may still be needed for forensic analysis in case there is a security incident involving the NNM server. Once the log files are configured correctly, if adverse changes are made to the NNM server the NNM administrator will be able to find out who was logged into the NNM server during the change. The NNM administrator should choose the option of overwriting events as the file size increases, which will drop one event at a time instead of an entire day all at once.

Compliance: This test is necessary to ensure the NNM administrator is aware of the security implications of allowing the log file defaults to remain unconfigured. The auditor will know whether the NNM server is in or out of compliance when using the testing steps outlined below.

Testing: The auditor can go to Start > Programs > Administrative Tools > Event Viewer, and then right-click on each log tab and select Properties. It should be verified that the maximum log size field has a value of at least 512 if not more. When the maximum log file size has been reached the overwriting option that should be chosen is the first option that will overwrite events as needed.

Objective/Subjective: This test produces objectively verifiable results that are outlined completely in the testing area of this audit step. The Windows log file configurations must match to be in compliance.

Step 27: Check NNM server for latest OS and IIS patch installation

Reference: http://v4.windowsupdate.microsoft.com/en/default.asp

Control Objective: Verify that the Windows 2000 operating system (OS) that the NNM system is running on has the latest security patches installed for the OS and the Internet Information Services (IIS) web server.

Risk: If the latest patches are not installed for the OS and the web server, the NNM server could fall victim to an attack that would allow someone to take control of the NNM server. An attacker could also launch a DoS attack, which would render the NNM system inoperable. Either attack could allow the network to go unmonitored for some time, data to be destroyed or stolen, or even used as a launching pad to attack many other devices that are in a trust relationship with the NNM server for monitoring purposes.

Compliance: The NNM server's OS and IIS server will or will not have been updated with the latest patches from Microsoft. If the Windows Update page lists patches that are needed, then the NNM server will have failed the test.

Testing: The auditor simply needs to open the NNM server's web browser and select Tools from the File menu and then Windows Update. Microsoft's Windows Update site will come up, and may prompt the auditor to trust plug-in data from Microsoft Corporation; if so, it is fine to click the Yes button. The link Scan for updates should be clicked on, and then wait for the site to search for critical updates that the NNM server needs to install. If the NNM server does not have one of the latest Critical Updates or Service Packs, the site will prompt for you to Review and install updates. Click this link, and read the description of each recommended update. Only critical updates for the NNM server will be listed, which are there to either repair discovered security vulnerabilities or eliminate bugs in the code to improve stability. Before any

30

updates have been installed, the NNM administrator needs to make sure that there is a good backup of the entire NNM server's hard drive. Once the auditor feels comfortable with the update recommendations, the Install Now button can be pressed. This will prompt the auditor to Accept the list of chosen updates before installation. Once Accept is pressed it will download the updates which will then prompt for the NNM server to be rebooted. The NNM administrator will need to find the next time the NNM server can be down for a reboot, and schedule it for a Change Control outage period. The auditor can then check the Automatic Updates, to ensure any future updates will be downloaded automatically and prompt the NNM administrator to install them. Navigate to Start > Settings > Control Panel > Automatic Updates, the checkbox at the top should be checked, and then choose the second radio button option which will automatically download critical updates in the future, but will wait for the NNM administrator to decide to install them on the NNM server. Once again, a full backup of the NNM server should be performed by the enterprise's system backup solution. The NNM server should use the same system patching policy as the enterprise's mission-critical Windows servers. Most corporate policies will be fairly similar to these patching steps and update configurations.

Objective/Subjective: This is an objective test that assures the NNM server remains current with OS and IIS patch levels. Although, the corporate policy will dictate whether the NNM administrator is able to comply with this objective.

Step 28: Checks the NNM server's hard drive maintenance routine

Reference: Personal Experience

Control Objective: These steps are to verify that the NNM server's hard drive is in good working condition for the security of the information contained within the NNM database and maintained by the NNM system.

Risk: If there were a problem with lack of disk space or a corrupted file system located on any partition within the NNM server, the NNM system could experience a period of unexpected downtime. NNM system downtime can cause a range of financial expenses, from SLA penalties for not monitoring a client's network, to increased labor needed to restore the NNM server from backups.

Compliance: The NNM server's hard drive maintenance schedule should conform to the enterprise's hard drive maintenance policy for mission-critical servers. The test results for this auditing step are binary.

Testing: The auditor should compare the enterprise's server maintenance policy to the hard drive maintenance schedule employed by the NNM administrator, which most likely will not be written. The free space on the local disk where the NNM installation resides should be checked to ensure there is enough free space for NNM system backups, NNM database growth, and NNM system and server logs. The NNM server should always have more free space than used space on the local hard disk. The auditor can double-click on My Computer and then select the local disk that the NNM installation resides on. This will show the current free/used space ratio on the NNM server. The NNM server should also have its local disks defragmented regularly, to ensure the security of the NNM system data and maximum NNM system uptime. The auditor should go to Start > Programs > Accessories > System Tools > Disk Defragmenter. The Disk Defragmenter utility will come up, and each disk should

31

be selected from the list and then the Analyze button should be pressed. The utility will analyze the disk to see if it requires defragmenting for optimal performance and data security. If the disk needs to be defragmented the utility will prompt for that action once the analyzation phase has ran. The disks should only be defragmented during the period of time that is set aside for Change Control tasks. This will ensure that while the OS is using a large amount of CPU time, the NNM system will not be expected to keep up the same level of network device monitoring, since it may fall behind in status polls.

Objective/Subjective: These are objective tests that will produce expected results. The procedure implemented by the NNM administrator may have to conform to the enterprise's mission-critical server maintenance schedules.

Step 29: Run a Nessus vulnerability scan against the NNM server

Reference: http://www.sans.org/rr/papers/5/78.pdf

Control Objective: This step is focused on a vulnerability scan against the OS and the IIS Web Server. This will ensure that anything that was missed during the auditing process, will be found by the scanner and corrected before an attacker discovers and exploits it.

Risk: If the NNM server has any vulnerabilities that were missed during the specific auditing steps, the free vulnerability scanner will find any that are left. If a vulnerability is found, Nessus will give basic steps to correct the problem, which should be given to the NNM administrator to correct as soon as possible. If these holes aren't corrected quickly, an attacker will find them and exploit them. This could cause the NNM server to be taken over, rendered unusable, or collected data destroyed. Vulnerabilities within the Windows OS and Microsoft's IIS Web Server are the most common vulnerabilities found on a Windows server and are often the easiest to exploit by an attacker. Nessus is a vulnerability scanner that is widely used by auditors to quickly run through several vulnerability tests that would take hours to do one test at a time. Nessus makes its way through a large list of plug-ins (vulnerability tests) that are written by many sources and submitted to the Nessus workgroup. The new plug-ins are made available to download, and then the most recent plug-ins are packaged with the next major Nessus release.

Compliance: This step will find many security hardening recommendations for the NNM server. The auditor will turn over the list of recommendations, which Nessus will create in its report, over to the NNM administrator who should correct them as soon as possible.

Testing: The Nessus architecture is broken into two different pieces. The Nessus server application must be run on a Linux computer, and the Nessus client can be ran on either Linux or Windows. Most auditors prefer to run the Nessus client from a Windows machine, because the Windows version of the client is much easier to use than the Linux version. Both Nessus pieces can be downloaded from http://www.nessus.org/download.html, and installed in minutes. Detailed steps of the installation options and setup of both pieces is beyond the scope of this paper, but the information needed for the setup can be found at http://www.nessus.org/demo/index.html. A Nessus scan should only be preformed during a change control period or during low network utilization. Permission for this

32

test to be completed must be given by an upper security manager or CIO, because of the attacking nature of the Nessus scan. Even if the DoS tests aren't ran against the NNM server, it could still cause the NNM server to behave unusually during the tests. An item that the auditor can expect to find in the results of a Nessus scan of a default Windows Server 2000 installation would be, for example, the capabilities to complete a null session with the OS. The items in the recommendation list of the Nessus report should be taken seriously by the NNM administrator, because a misconfiguration or vulnerability that seems simple can be exploited by an attacker just as simply as it would be for the recommendation to be corrected by the NNM administrator.

Objective/Subjective: The results of this test are very objective. The list of recommendations that are created in the Nessus report are described thoroughly and steps to correct the vulnerabilities are mentioned. The auditor may be asked to fix the problems that are found, but this may not be a part of the scope that was defined for the auditor before testing began.

Assignment 3 – Audit Evidence

3.1 – Conduct the Audit

Each step discussed in the top 10 list below, of most important security tests for an NNM system, can be seen in their entirety in the audit check list of 29 items above. This section of the paper calls for the results of the audit to be discussed concerning only the top 10 most important auditing steps from the checklist, including visual output documentation and overall assessments of each test. Only the testing section of each step was included in this section to avoid redundancy.

Some of the steps that were used to audit the live NNM system may have test results that the NNM administrator was fully aware of, but could do nothing about because of the enterprise's corporate policy or other environment specific limitations. These have been discussed when needed in the full audit checklist above and the audit report below.

The steps that were chosen from the full audit checklist to be used in the list of top 10 security tests were picked, because of their importance in an environment utilizing and depending on an NNM system. Other large security risks were discussed in the full audit checklist, but this top 10 list of NNM security risk audit steps was formed to cover major NNM specific items that an NNM auditor would need to verify compliance.

There are many best-practice and auditing papers that are written concerning securing and hardening Windows OSs and Microsoft IIS web servers. An attacker would probably check vulnerabilities within these two options first, but the focus on these 10 steps will be on much needed SNMP and NNM system checks. Microsoft

33

specific topics are included in this paper, because this NNM system is built on and around Microsoft software that it depends on.
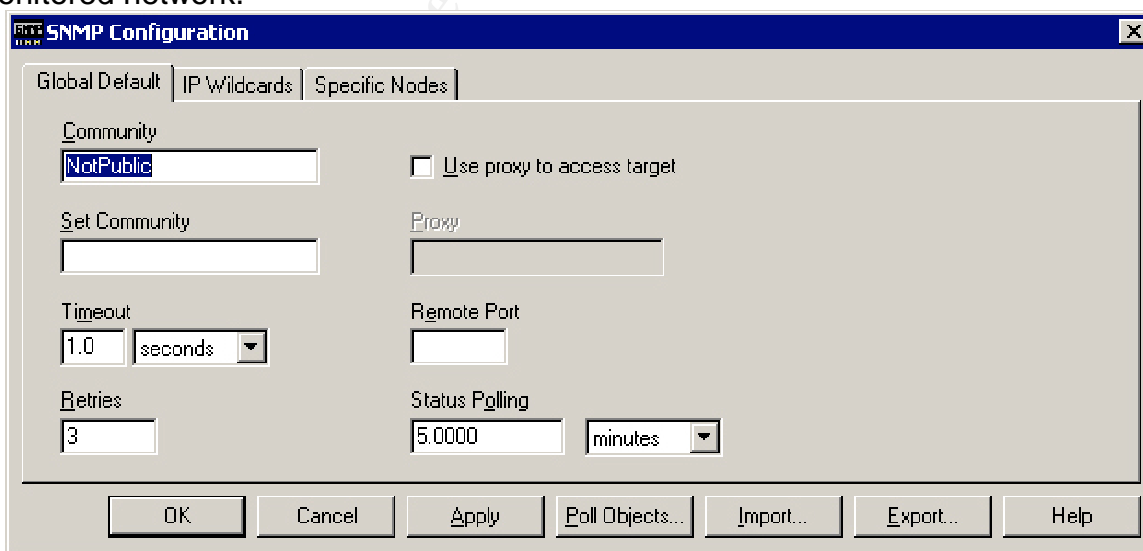
NNM remote control issues and the combination of a network's SNMP security posture is the focus of this audit, and that is currently missing from the IT Security community's available resources.

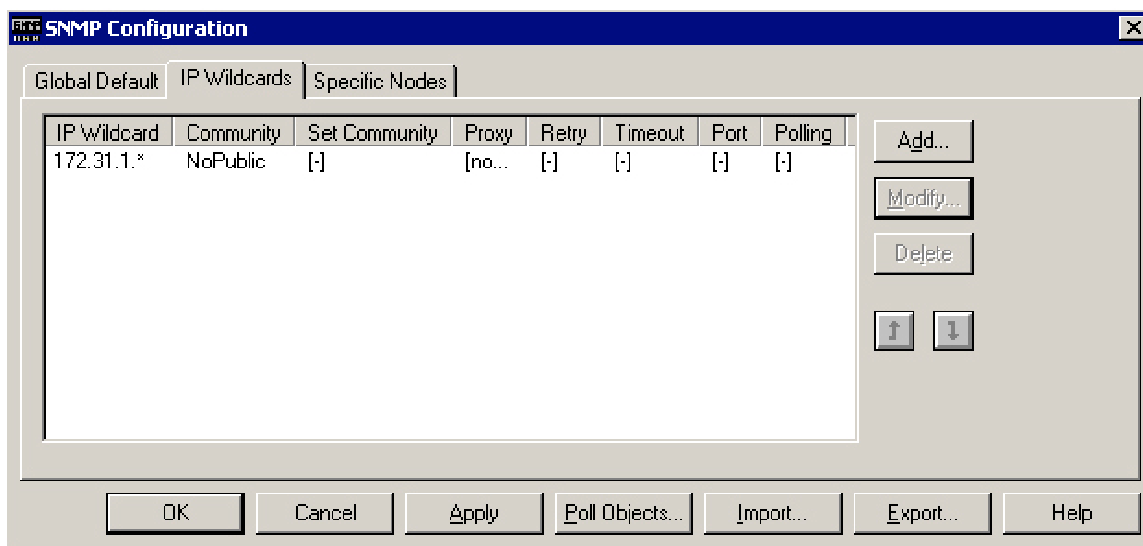*10 Highest Security Concerns of an NNM System from the Audit Checklist*

1. Step 3: *Public* should be removed as the SNMP RO string from all devices
2. Step 4: SNMP RW string should be removed from all SNMP-enabled devices
3. Step 8: Check that NNM root directory Share specifies authorized NNM users
4. Step 9: Lock down web access to specific authorized NNM users
5. Step 15: Verify the NNM management station is running the latest NNM release
6. Step 16: Verify NNM backups run locally, weekly, and are stored remotely
7. Step 17: Check NNM patch level for compliance with most recent release
8. Step 21: Check for commonly overlooked NNM performance issues
9. Step 24: Check to see who has administrative privileges on the NNM server
10. Step 27: Check NNM server for latest OS and IIS patch installation

Step 3: *Public* should be removed as the SNMP RO string from all devices

Testing: On the NNM server, the RW map can be opened and the auditor can then go to Options > SNMP Configuration in the File menu. Under the Global Default tab, there is a text field that contains the RO community string; this should be something other than *Public*. The Set Community text field below should be empty for an NNM monitored network.

Under the IP Wildcards tab, if there are networks explicitly defined to use a Community string other than the default global RO string, the devices on this network should be changed to use something other than *Public*.

34

Under the Specific Nodes tab, if there are nodes explicitly defined to use a Community string other than the default global RO string, the device should be changed to use something other than *Public*.

Results: It was verified that the system complied with 2 of the 3 checks outlined in the testing area of this audit step. The NNM system used an SNMP RO string other than *Public* in all three areas that were to be checked by the auditor other than the Specific Nodes section of the SNMP Configuration for the NNM system. This area caught many devices that are still configured with the default RO string of *Public*. This should be changed immediately.

Passed/Failed: The NNM system failed the test of having all areas of the SNMP Configuration using an SNMP RO string other than *Public*.

Step 4: SNMP RW string should be removed from all SNMP-enabled devices

Testing: The auditor should open the RW map and navigate to Options in the File menu and then select SNMP Configuration. This will open the SNMP Configuration window, where under each tab the auditor can verify that there aren't any RW strings defined globally, by subnet, or by specific devices.
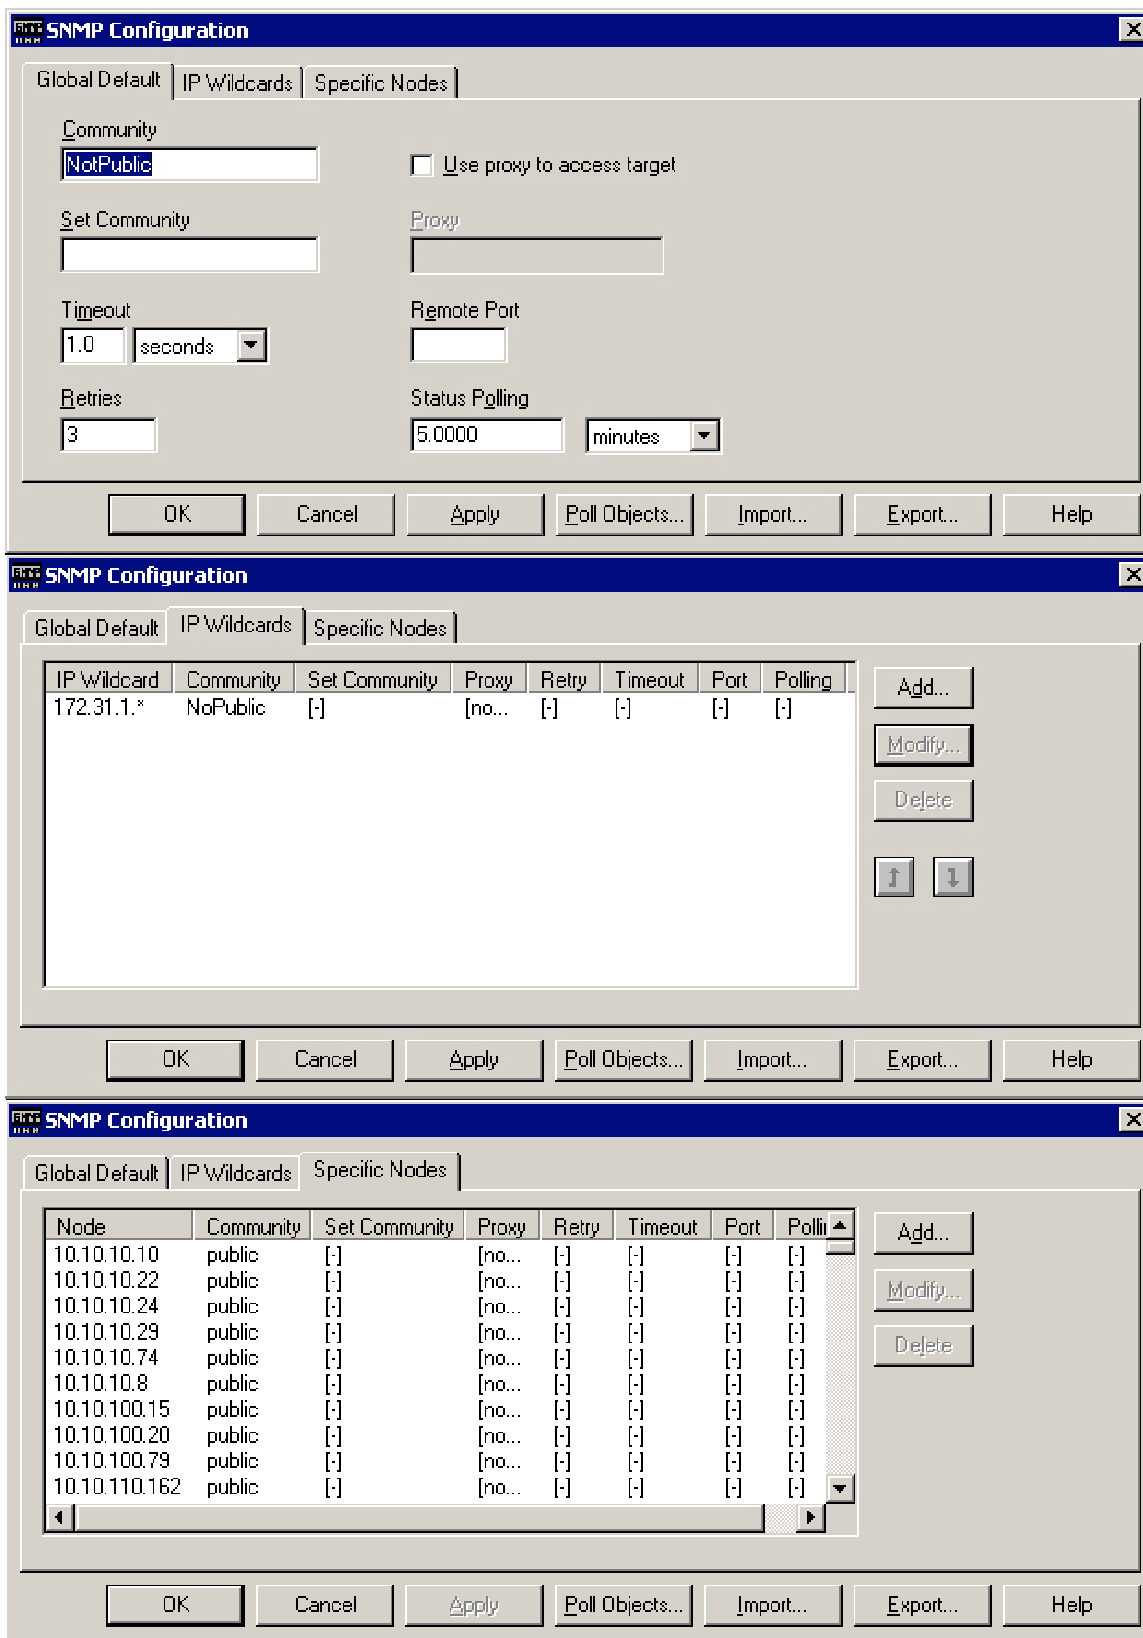
The auditor should also review the enterprise's verbal or written SNMP policy, and verify that it explicitly defines that all SNMP-enabled devices should be stripped of
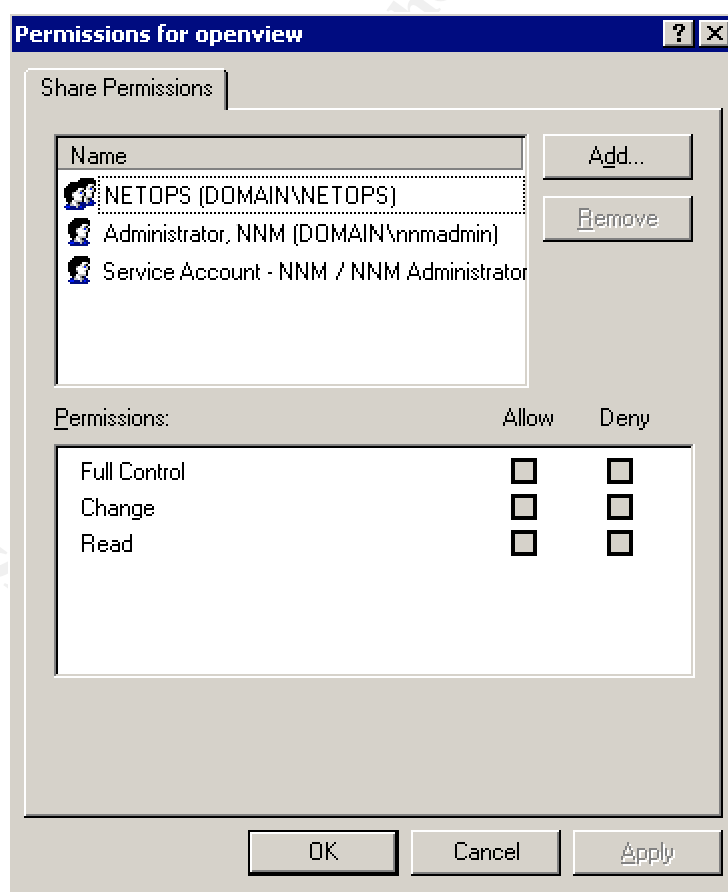
their RW string before they are placed on the network.

Results: All three sections under the SNMP Configuration window were checked for the use of an SNMP RW string, and there weren't any found. The NNM system is not using any RW strings, which is important for an NNM system that is only being used for network device monitoring. The NNM administrator was asked about viewing the corporate SNMP policy. There were no formal documents to review, but the NNM administrator assured me that he and the Network Security Team manager strongly agree in having all default RO strings updated and RW strings stripped from all devices that are placed on the network. They regularly discuss the importance of the SNMP configuration with the managers of other teams that are in charge of certain devices that are placed on the network with SNMP capabilities enabled.

Passed/Failed: The NNM system passed the RW string configuration test. The NNM system passed the subjective verbal SNMP policy test, but a written SNMP policy is needed. The documented corporate policy would dictate its own self-regulation of compliance.

Step 8: Check that NNM root directory Share specifies authorized NNM users

Testing: The NNM installation path would need to be checked for compliance to the item described. Right-clicking, selecting Sharing, and then the Permissions button on the root installation folder of NNM, will show which users have been setup with file access.

The NNM administrator would need to make a list of users and their privilege level in

this folder, and then the auditor can check for compliance. The administrator will need to be reminded that this configuration not only effects whether or not a user will have access to the NNM .conf files, database, and topology files, but also that the configuration decides what level of control the NNM remote management clients have over the NNM maps.

Results: The list of authorized NNM users given to me by the NNM administrator was used when checking who had Administrative privileges on the NNM installation path. The test checked out correctly, with only the NNM administrator having Full Control. The Service Account had Change-level control, which was needed to run the *ovbackup.ovpl* script during NNM system backup time periods. The NetOps group had Read-level control, which enabled those users to use the NNM maps for monitoring purposes.

Passed/Failed: The NNM server passed this test with full compliance.

Step 9: Lock down web access to specific authorized NNM users

Testing: Look for *$NNMserver\www\etc\htpasswd* to see if the NNM administrator has modified the User Authentication Password File with authorized NNM user names.
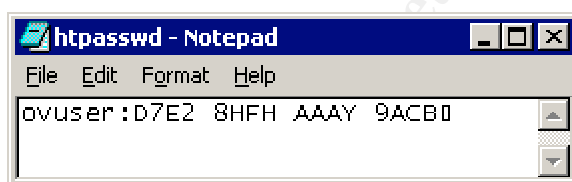
The auditor should then point a web browser to http://$NNMserver/OvCgi/ovlaunch.exe and see if they can connect. The auditor should not be able to connect without a valid username and password.
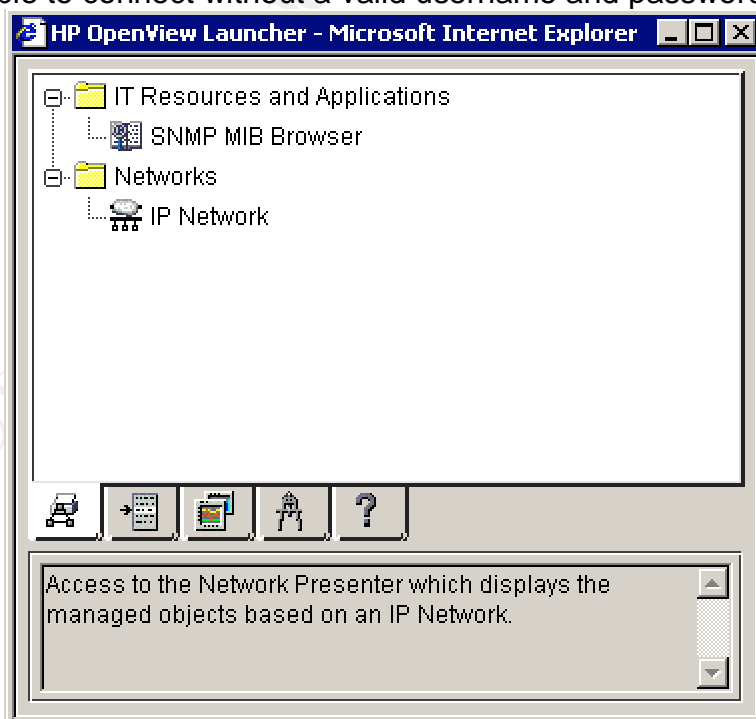
Results: The NNM system was not configured to use a username and password

when logging into the NNM web map and SNMP tools. The default *ovuser* account was the only one found in the *htpasswd* text file. I was able to point a web browser at the NNM web launcher from another machine on the network and gain full access to the NNM web maps and NNM SNMP tools, without being prompted for a username and password.

Passed/Failed: The NNM system failed this test, and should be corrected immediately using the steps outlined in this test.

Step 15: Verify that management station is running latest major NNM release

Testing: The auditor can simply point a browser at http://openview.hp.com/products/nnm/index.html to check for any new major releases that are available.

A new release upgrade can only be accomplished if the enterprise has a software subscription license for NNM with HP. The management station map can be pulled up to compare the current release version to what is running on the management station. Once the map is pulled up the auditor can navigate to the Help button in the File menu and then choose About HP OpenView from the drop down list.

39

The second line on this window will state the current version of the NNM system at the end of the line: HP OpenView Windows NNM Release. The X's in the string (0X.XY) compose the major version release number.

Results: The new NNM version 7.0 has been released, but the NNM system is still running at version 6.4. The enterprise has a current software subscription for NNM through HP, but the NNM administrator said that he wasn't comfortable upgrading the production NNM system until the first patch has been released for the newest NNM version to repair bugs in the new code.

Passed/Failed: The NNM system failed to pass the test, but the need to jump from a stable and fairly recent major NNM release to the latest and greatest isn't entirely necessarily in this situation.

Step 16: Verify NNM backups run locally, weekly, and are stored remotely

Testing: A discussion should first take place between the auditor and the enterprise's backup solution system administrator. The enterprise's written backup policy should be reviewed to ensure that the NNM system is being backed up at least once a week. The backup policy should also define a reasonable disaster recovery plan. The backed up NNM system data should be offsite in a different physical location of the enterprise's campus. The time of the NNM system backup can then be determined by checking the IP or DNS name that is set in the backup system's configuration. The NNM server data backup should happen directly after the *ovbackup.ovpl* script has run within the NNM system. The auditor should log into the NNM server and verify that there is more free space than used space left on the hard drive, which will ensure NNM server stability after backups have been run locally of the critical NNM system configurations.

40

The Scheduled Tasks window should then be pulled up by going to Start > Settings > Control Panel > Schedule Tasks. There should be a scheduled task that runs the *ovbackup.ovpl* script at least once a week, and directly before the enterprise's backup solution does the weekly backup of the NNM server's hard drive.

Results: The NNM server is being backed up by the enterprise's backup solution once a day. The disaster recovery plan is exceptional, accounting for daily backups on magnetic tapes, which reside in a different area on campus from the NNM server itself. The tapes are then copied and taken off campus to an entirely different geographic region where they are stored for several months. The enterprise's backup solution does not run directly after the NNM system backup scripts run. The NNM server hard drive that houses the full NNM installation, does not have more free space than it does used space. The *ovbackup.ovpl* script does run once a week, but not directly before the full NNM server backup.
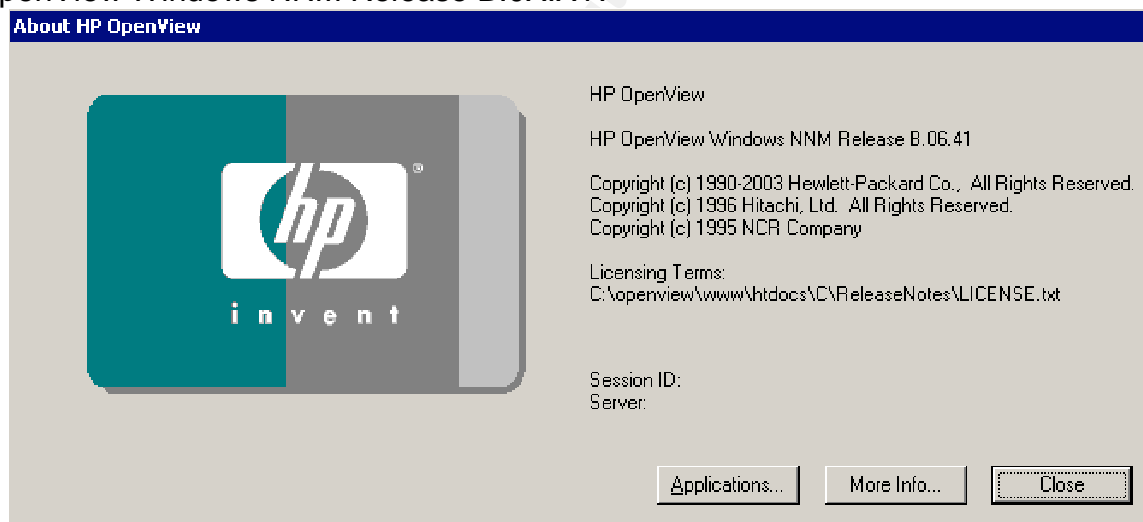
Passed/Failed: The NNM server passed the enterprise backup solution test, and the disaster recovery test. The NNM server failed the enterprise backup solution timing test, but because the NNM server is backed up daily and the NNM system is only backed up weekly this was given a grade of passing. The NNM server failed the hard drive free space check, but because the NNM server has multiple GB of free disk space this was also given a grade of passing.

Step 17: Check NNM patch level for compliance with most recent release

Testing: The auditor can log onto the NNM server and open the RW map. While the map is being loaded, the splash screen will show the current consolidation patch level of the system. This screen can also be brought up by navigating to Help in the File menu and then selecting About HP OpenView. The second line of text will state HP OpenView Windows NNM Release B.0X.XY.

The X's are the major release version number, and the Y value is the current consolidated patch level of the system. There are other patches that are released that will not affect the Y value when applied to the system. Most of these patches are quick fixes for a discovered bug or vulnerability, and after enough of the intermediate patches are released there is a consolidated patch that rolls up several new fixes, but also holds all of the old ones as well. These affect the value that is represented by the Y, and all subsequent patches that are released will be functionally dependent upon the prior consolidated patch and can't be installed without them. The site http://support.openview.hp.com/cpe/patches/nnm/6.4x/win.jsp should be referenced to find out the latest consolidated patch release number.
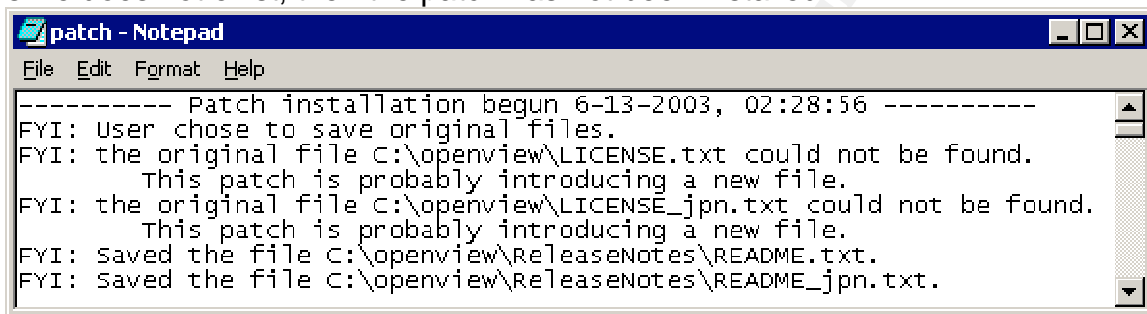
NNM_00998

Description:	Consolidated Patch 1
Posted Date:	2003-May-29
Symptoms:	**Cumulative Consolidated Patch** NNM_00988: 8606284920: E-mails generated by reporting are not received, if qmail is configured as the mail server. 8606284812: On Remote Console, xnmevents takes a long time to come up. 8606207670: On configuration updates, snmpCollect does not reuse existing collections. 8606228571: The migratable IP address reported by Sun Clusters are not considered as migratable by NNM. 8606288064: netmon exits on receiving an invalid PDU. *...Symptom text truncated. Please refer to the patch text file for a complete list of symptoms.*

To test whether or not the latest consolidated patch has been applied to the NNM system, the auditor can navigate to $NNMserver\Patches\Patch_Name\patch.txt. If this file does not exist, then the patch has not been installed.

```
---------- Patch installation begun 6-13-2003, 02:28:56 ----------
FYI: User chose to save original files.
FYI: the original file C:\openview\LICENSE.txt could not be found.
        This patch is probably introducing a new file.
FYI: the original file C:\openview\LICENSE_jpn.txt could not be found.
        This patch is probably introducing a new file.
FYI: Saved the file C:\openview\ReleaseNotes\README.txt.
FYI: Saved the file C:\openview\ReleaseNotes\README_jpn.txt.
```

The site http://support.openview.hp.com/cpe/patches/nnm/6.4x/win.jsp should be referenced to find out the latest intermediate patch release number.
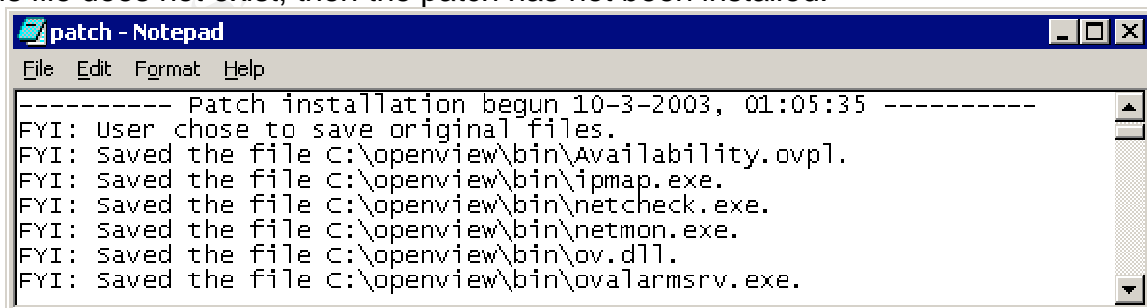
NNM_01008 | This patch depends on: NNM_00998

Description:	Patch for Sep 2003
Posted Date:	2003-Oct-15
Symptoms:	Change Request: 8606311794 snmpCollect does multiple collections on same node-instance combination. Change Request: 8606312494 ovdwevent does not work properly on bad PDU. Change Request: 8606325594 'owwrs' does not start on Solaris. Change Request: 8606322563 Pop-up messages are shown in Web Alarm Viewer, whenever there is a change in the trapd.conf file through 'xnmtrap' utility. *...Symptom text truncated. Please refer to the patch text file for a complete list of symptoms.*

To test whether or not the latest intermediate patches have been applied to the NNM system, the auditor can navigate to $NNMserver\Patches\Patch_Name\patch.txt. If this file does not exist, then the patch has not been installed.

```
---------- Patch installation begun 10-3-2003, 01:05:35 ----------
FYI: User chose to save original files.
FYI: Saved the file C:\openview\bin\Availability.ovpl.
FYI: Saved the file C:\openview\bin\ipmap.exe.
FYI: Saved the file C:\openview\bin\netcheck.exe.
FYI: Saved the file C:\openview\bin\netmon.exe.
FYI: Saved the file C:\openview\bin\ov.dll.
FYI: Saved the file C:\openview\bin\ovalarmsrv.exe.
```

Results: The NNM system had the most currently released consolidated patch installed on the system. The NNM system did not have the most currently released

intermediate patch installed on the system. The NNM website showed the latest intermediate patch was NNM_01008, and the NNM system was running NNM_01006.

Passed/Failed: The NNM system passed the consolidated patch test, because NNM_00998 was already installed. The NNM system failed the intermediate patch test, because the last intermediate patch to be installed was NNM_01006, while NNM_01008 had been released.

Step 21: Check for commonly overlooked NNM performance issues

Testing: The first test will be for CPU load. This can be accomplished by logging onto the NNM server, hitting *CTL+ALT+Del*, and then selecting Task Manager. When the Task Manager window comes up, the Performance Tab should be selected and the current CPU load should be monitored for at least a minute to verify that the CPU never reaches a 100% load.

The performance of the NNM system's primary purpose can be measured by a graphing tool built into the maps. The auditor can open the NNM RW map, select Performance from the File menu, and then select Network Polling Statistics. The subsequent window will have a real-time graph that shows the status polls and SNMP polls that are running in the background to the monitored devices.

All of these graphs should stay above zero the majority of the time; otherwise, other network polling configurations may need to be looked at to be reconfigured. Both the Task Manager and the Network Polling Statistics graph should be left open and watched by the auditor to verify that the graphs stay in the acceptable ranges decided upon in the previously discussed steps. The auditor should then open a command prompt window on the NNM server, and type in the command *ovstatus –v*. This will list the NNM processes that should be running. The list should be scanned to verify all processes are up and running.

If one of the processes is not running for some reason, the command *ovstart –c*

service (the name of the service not running). The command *ovstatus –v netmon* should then be ran in the command line. The output we are concerned with here is whether or not there are overdue polls and how far behind they are.

```
Command Prompt                                                    _ □ ×

Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.

C:\WINNT>ovstatus -v netmon
 object manager name: netmon
 state:                  RUNNING
 PID:                    7592
 last message:           Initialization complete.
 exit status:            -
 additional info:
16:52:36 Polling 949 interfaces, 11388 polls/hour.  6 overdue polls, current max
imum 2 seconds behind.  Worse was 6 polls a maximum of 2 seconds behind at 10/15
/03 16:52:36.  100% interfaces available.


C:\WINNT>_
```

   As before, if the NNM system is falling behind, certain Network Polling parameters may need to be reconfigured to allow the network to be adequately monitored. Another concern is whether or not the NNM system is having problems doing name resolution for devices on the maps. If this command gives us output that details multiple Domain Name Services (DNS) requests with high response time, the NNM administrator may need to setup a local DNS server on the NNM server. The NIC would then use the local DNS server as the primary and the enterprise's primary DNS server as the NNM server's secondary DNS server. A final step that is sometimes overlooked is that of the speed of each device when answering the NNM system during an SNMP configuration check. If a device makes the NNM system wait to complete a scheduled poll, it will be during a full SNMP configuration check. If a device stalls during an SNMP configuration check it is usually during the routing table call. Routing tables are called to aid in IP device discovery and topology configuration changes that should be reflected on the NNM maps. This will only be a problem on routers with large routing tables, which are possibly connected to the network with slow links, and little CPU to spare. Border routers are usually the culprit, and can take many minutes and sometimes hours to dump out their routing tables to the NNM system. Because border routers are at the edge of a network these routing tables aren't always that important for IP device discovery on the internal network. These devices should be selected by the auditor on the RW map, Fault should be selected on the File menu, Network Connectivity, and then Poll Node. This will open the SNMP polling window which will run through an ordinary poll step by step, and the system will probably stall when trying to retrieve the routing table.

46

Poll Node

File   View   Help

Name or address:

border.router.net

Poll Results

```
16:57:46    Get interface ifAlias'
16:57:46    Get IP forwarding status
16:57:46    Get system location
16:57:46    Get system contact
16:57:46    Get default routing entry
16:57:46       SNMP error: no such name
16:57:46    Get routing table
```

Stop   Restart   Close

This can cause unneeded stress on the border routers, and cause the NNM system to become behind in its polling device list. The NNM administrator may want to consider specifying within the SNMP Configuration area that this device does not need to be a part of a full SNMP configuration check.

Results: The Task Manager CPU graphs on the NNM server never reached and remained at 100% CPU for more than a short period of time. All graph lines in the Network Polling Statistics stayed above zero except the line labeled Seconds Until Next SNMP Poll. All of the major NNM processes were running when tested. The *netmon* service was running steadily without falling behind by more than a few seconds. A secondary DNS server was running on the NNM server to help alleviate the DNS lookup stress. The border routers did stall during the SNMP poll, which may be causing unneeded stress on the NNM system or possible delays in the polling of other devices next in line.

Passed/Failed: The NNM server passed the Task Manager test. The NNM system failed the Network Polling Statistics test, because the Seconds Until Next SNMP Poll line in the graph fell below zero. This test result can be up for interpretation by the auditor, but the test calls for all lines to be above or equal to zero. The NNM system passed the process list test. The NNM system passed the *netmon* service test. The NNM server passed the secondary DNS test. The NNM system failed the border router test, because it stalled while retrieving the routing table during the SNMP configuration poll.

Step 24: Check to see who has administrative privileges on the NNM server

Testing: The auditor can use the list of authorized NNM users and verify using Microsoft Computer Management that only NNM administrators are also Windows administrators on the NNM server.

The auditor will then try to connect remotely using Microsoft Terminal Services or Computer Management and see if a user not specified as an NNM administrator can connect to the NNM server.

The auditor will also need to determine whether an unauthorized NNM user account is able to connect to the NNM installation path using an NNM remote management client and is able to open an NNM map.

Results: The list of authorized NNM users that was given to me by the NNM administrator matched what was found in the NNM server's list of OS administrators. I was unable to connect to the NNM server using Terminal Services and Computer Management when logged onto a different machine with a username other than an OS administrator. I was also unable to connect to the root NNM installation path when trying to open an NNM map using a remote NNM management client on a different machine using a username that is not an OS administrator.

Passed/Failed: The NNM server passed the OS administrator test. The NNM server also passed the Terminal Services and Computer Management test. The NNM system passed the NNM remote management test.

Step 27: Check NNM server for latest OS and IIS patch installation

Testing: The auditor simply needs to open the NNM server's web browser and select Tools from the File menu and then Windows Update. Microsoft's Windows Update site will come up, and may prompt the auditor to trust plug-in data from Microsoft Corporation; if so, it is fine to click the Yes button. The link Scan for updates should be clicked on, and then wait for the site to search for critical updates that the NNM server needs to install. If the NNM server does not have one of the latest Critical Updates or Service Packs, the site will prompt for you to Review and install updates.

49

Windows Update

Home | Windows Catalog | Windows Family | Office Update | Windows Update Worldwide

All Products | Support | Search | microsoft.com Guide

Microsoft

Windows Update

- [ ] Welcome
- [ ] Pick updates to install
  - [ ] Critical Updates and Service Packs (4)
  - [ ] Windows 2000 (8)
  - [ ] Driver Updates (1)
  - [ ] Review and install updates (4)

Other Options
- [ ] View installation history
- [ ] Personalize Windows Update
- [ ] Get help and support

See Also

Total Selected Updates

Review and install your selected updates

Your total selected updates include an exclusive item that must be installed separately from other updates. To install the exclusive item, click **Install Now**. If you wish to install other critical updates, remove the exclusive item by clicking **Remove**.

[Install Now]   Total: 4 = 7.6 MB, 1 minute

Windows 2000 Service Pack 4 Express Install for End Users*
Download size: **589 KB**, < 1 minute
Download Windows 2000 Service Pack 4 to obtain updates that can help improve security, application compatibility, operating system reliability, and Setup. This service pack also includes the updates contained in previous Windows 2000 Service Packs. Read more...

* Must be installed separately from other updates

[Remove]

October 2003, Cumulative Patch for Internet Explorer 6 Service Pack 1 (KB828750)
Download size: **2.1 MB**, < 1 minute
Security issues identified in Microsoft Internet Explorer (IE) could allow an attacker to compromise systems with IE installed (even if IE is not used as the Web browser). For example, an attacker

Click this link, and read the description of each recommended update. Only critical updates for the NNM server will be listed, which are there to either repair discovered security vulnerabilities or eliminate bugs in the code to improve stability. Before any updates have been installed, the NNM administrator needs to make sure that there is a good backup of the entire NNM server's hard drive. Once the auditor feels comfortable with the update recommendations, the Install Now button can be pressed. This will prompt the auditor to Accept the list of chosen updates before installation. Once Accept is pressed it will download the updates which will then prompt for the NNM server to be rebooted. The NNM administrator will need to find the next time the NNM server can be down for a reboot, and schedule it for a Change Control outage period. The auditor can then check the Automatic Updates, to ensure any future updates will be downloaded automatically and prompt the NNM administrator to install them. Navigate to Start > Settings > Control Panel > Automatic Updates, the checkbox at the top should be checked, and then choose the second radio button option which will automatically download critical updates in the future, but will wait for the NNM administrator to decide to install them on the NNM server.

Automatic Updates

Windows can find the updates you need and deliver them directly to your computer.

☑ Keep my computer up to date. With this setting enabled, Windows Update software may be automatically updated prior to applying any other updates.

Learn more about automatic updating

Settings

○ Notify me before downloading any updates and notify me again before installing them on my computer

⦿ Download the updates automatically and notify me when they are ready to be installed

○ Automatically download the updates, and install them on the schedule that I specify

[Every day ▾] at [3:00 AM ▾]

Learn more about scheduled installing

[Declined Updates]

[OK]   [Cancel]   [Apply]

Once again, a full backup of the NNM server should be performed by the enterprise's system backup solution. The NNM server should use the same system patching policy as the enterprise's mission-critical Windows servers. Most corporate policies will be fairly similar to these patching steps and update configurations.

Results: The NNM server did not have the most recent patches and OS service packs installed. This may or may not be in line with the corporate policy, but nevertheless is not best-practice. The NNM server did have Automatic Updates configured correctly.

Passed/Failed: The NNM server failed the OS and IIS web server patching test. These patches and service packs should be installed as soon as possible. The NNM server passed the Automatic Updates test.

3.2 - Measure Residual Risk

A great deal of risk is associated with an NNM system running on Windows Server 2000 with Microsoft IIS web server. This particular NNM system has an unnecessary amount of risk still associated with it. The NNM administrator has taken several steps to help protect the enterprise's valuable network monitoring system, but he must remain aggressively proactive in his fight to secure the valuable data and service that the NNM system gives to the company.

The results from the audit carried with them a certain amount of risk, but for the

51

most part nearly all of the tests that failed can be corrected on the NNM system and server. A few of the test results that failed had fallen victim to circumstance, such as tests concerning patching the OS, the NNM system, and keeping up with the latest NNM releases. The NNM administrator usually installs the latest NNM system patches during the next change control period after new patches are released. He is on the mailing list for all NNM system updates, patches, and major releases. That particular auditing step happened to be run the day of a new patch release, and the next change control period was not for another day.

The latest release for NNM, version 7.0, could be installed and licensed using the company's software subscription with HP, but the NNM administrator feels more comfortable waiting for the first patch to be released by HP on the new version before he considers installing the latest version. There are two items of risk that must be considered here: it is risky for the NNM administrator to install the latest version that may not be as stable as the current version, but there is also a risk in running a version that may have more vulnerabilities that are now corrected in the latest NNM version. These items must be discussed, or a corporate version upgrade policy on mission-critical systems may also dictate how the NNM administrator must proceed in this kind of a situation.

The NNM server's OS is also at risk, because the latest vulnerability patches and service packs have not yet been installed on the NNM server. This cannot always be corrected quickly, because the corporate policy calls for all mission-critical servers to have service packs tested before being applied to a Windows server. These precautions are in place to minimize the risk of downtime from a broken OS. The NNM administrator would also have to make sure there is a full backup of the NNM server before any major patches or service packs are installed on the NNM server.

Ideally the enterprise's backup solution would backup the NNM server directly after the NNM system backs itself up once a week. The NNM server receives a full hard drive backup once a day, and this lessens the risk of not having a known good-state copy of the NNM data and configurations coming directly after an NNM system backup. The timing of the NNM server backup is out of the hands of the NNM administrator, since the enterprise's backup solution runs through its list of mission-critical servers at a different time each day.

I would recommend for each test result, that options would be considered and brought forward to management on how to lower the associated risks. If the budget allows, the NNM administrator should be given a test server with adequate hardware to run each new major release of NNM in parallel to the production server, until he feels comfortable with a full cutover to the latest version. This will allow the NNM system to be using a more stable and secure version of NNM, and management will be pleased to have the added functionality that a new version of NNM may bring to the enterprise.

This testing of new NNM versions is very important, because the only *Public*ly

documented vulnerabilities that I could find described by CERT or other communities were concerning systems prior to NNM version 6.4. Keeping up with the latest releases from HP, will keep the risk of documented vulnerabilities away from the NNM system that is relied on so heavily by the enterprise.

There aren't a lot of major items that need to be done on the NNM system and server that would call for more hardware or manpower. Cost will remain low, but the NNM administrator will need to be more aggressively proactive, as mentioned before. A few security configurations on the remote management capabilities of the NNM system will take care of the majority of the problems that were found. Keeping up with NNM server OS patching is very important, and should be performed during a change control period after backup precautions have been taken.

The recommendation that would take the most time to accomplish would be setting up another NNM server with the newest release of NNM installed for testing. But this step would only take a day and would be very beneficial to the security and stability of the NNM system. The NNM server could be built in the morning, and the NNM system could be installed and configured in the afternoon. The NNM system will be able to discover the majority of the network in one night. The following day the NNM administrator would have the newest NNM release running parallel to the production NNM system. This new NNM system should be monitored for a week or more, before an upgrade of the existing NNM system should be considered.

All of these auditing steps make good business sense to follow, because the cost of time, money, and resources needed to accomplish these steps is far lower than the cost incurred by an organization with a compromised or extremely vulnerable NNM system that is mission-critical to the enterprise. The majority of the cost incurred will be that of labor as the NNM administrator sets up username and passwords for NNM web map access for each authorized NNM user. The NNM server's remote and local access to specific users will also need to be reviewed and updated periodically.

The majority of the control objectives were achieved successfully during the NNM system audit. The NNM system was using an SNMP RO string other than *Public* globally, but was allowing several devices on the monitored network to remain using the default RO string of *Public*. These devices were primarily Windows 2000 servers on the network. A server team member who is in charge of maintaining those machines should correct this immediately. All SNMP RW strings had been removed from the configuration area of the NNM system, and no specific nodes showed up as being allowed to use any RW strings throughout the monitored network.

The NNM installation directory on the NNM server was locked down to specific credentials, which kept unauthorized users from connecting to the NNM system with an NNM remote management client. It was also verified that only authorized NNM users had administrative privilege level access on the NNM server's OS. The NNM web access had not been configured to keep unauthorized users from pulling up an NNM web map.

The NNM management station was not running the latest major NNM release, which can be very important for many reasons; options on how to achieve this objective should be quickly addressed. The NNM system was missing an NNM intermediate patch, but out of the seven released, this was the only one missing and was scheduled to be installed during the next change control period. The OS was missing several patches and also a service pack, which had been overlooked by the NNM administrator. The NNM system was backed up weekly, and the entire NNM server's primary hard drive was copied off remotely on a daily basis.

The NNM system was configured well concerning NNM system settings that can cause performance issues for the NNM server. The NNM system didn't fall behind for more than a few seconds at a time, and the local DNS server corrected much of the name lookup latency. The CPU did reach 100% from every once in a while, but overall the NNM server was running well. The NNM administrator should focus on completing all of the NNM remote access objectives that were outlined in the audit steps. The OS should be patched constantly, because an OS vulnerability is one of the easiest and most likely attacks that would be exploited on an NNM server.

3.3 - Is the System Auditable?

Security hardening steps and best-practice NNM system configurations are auditable on a live NNM system. There are some auditing steps that were discussed that could be done on a NNM test system, but shouldn't be done on a live, mission-critical monitoring system. Although, some of the tests wouldn't validate the test results gained, unless the NNM system were setup to be enduring the same kind of stress that a live NNM system would feel on an aggressively monitored network.

The audit step I am talking about would be a network vulnerability scanner test using Nessus, which is discussed in the full audit checklist. If certain tests were performed by a Nessus session against a live NNM system, it could overload the system and crash the NNM server or keep it from performing its primary purpose of accurately monitoring the status of mission-critical devices on the network. The problem is that only a fully functioning and adequately taxed NNM system on a comparably busy network would accurately depict what would happen if an attacker were to run this type of vulnerability test against a live NNM system.

Other than this kind of extreme testing, an NNM system is definitely auditable. I only mention this testing to be extreme, because this is somewhat outside the scope of an NNM system audit. Although, this is one of the most important tests that can be run against a device, it is also risky to the stability of the NNM server and accuracy of the NNM system. This test should only be performed after a full NNM server backup, and during a change control period.

The NNM security checks that are discussed during the majority of the audit, are

54

the most important to focus on, because of the lack of this kind of documentation for an NNM system. There are some configuration settings that should be setup correctly on an NNM system that are often overlooked by an NNM administrator. The network-wide SNMP configuration is also a big part of a successfully secure and hardened NNM system. This can be difficult to audit because of the broad scope of the necessary tests, and the need to discuss so many far reaching changes with many different personnel groups and the system administrators themselves.

Assignment 4 - Audit Report

4.1 - Executive Summary

The NNM system and server audit was a success. The primary objectives were all checked for compliance against the NNM system and the NNM server. Some NNM system configuration settings will need to be changed, but they will be easy for the NNM administrator to implement. The performance of the NNM server is good. The SNMP configurations of the network devices that are being monitored are not setup as well as they could be. Recommendations will be made later concerning what can be done to further secure the NNM server and system, and also help mitigate the risk to the network devices that have less than optimal SNMP configurations.

The NNM system fully passed four out of the ten primary objectives. These 10 objectives are tested, because of their importance in reaching a fully hardened NNM system. These results do not mean that the NNM system is not stable or performing well, but it does mean that there is room for improvement if the NNM system is to be configured as securely as it can be.

This does require more time on the NNM administrator's part though; which is not to say that the NNM system has been neglected, but to reach absolute perfection, the system must be aggressively maintained. New patches for NNM, Windows, and the IIS web server come out weekly sometimes, and they must be applied as soon as possible to mitigate vulnerability risks that they are correcting within the NNM system and server. One of the biggest concerns, that can also be the most difficult to correct, is that of configuring each SNMP-enabled device on the network to be as secure as it can be from the protocols perspective.

The SNMP protocol is inherently insecure, but also very valuable. If the SNMP protocol is not used, then insight is lost into the current state of the monitored network of devices. Most system administrators of mission-critical devices depend on the SNMP messages to allow them to be constantly aware of the changes within the system. Because the NNM system is directly tied to the SNMP configuration of each device on the monitored network, these issues must be addressed within this audit.

A formal SNMP policy must be created and maintained by a security professional within the organization, which will state the importance of each device's compliance to the SNMP security configuration policy. This will take a lot of stress and time away from the NNM administrator's duties, because it would no longer be up to him to convince system administrators to take the necessary SNMP configuration precautions when securing their systems before deployment.

Overall, the NNM system is functioning well, and is reasonably secure. The NNM administrator has done a good job in trying to get the rest of the devices on the network securely configured in respect to the SNMP protocol. The OS's patch level is fairly up to date. The remote management capabilities of the NNM system has been locked down somewhat, so that only authorized NNM users can have any sort of control over the NNM maps and topology information. The hardware is keeping up with the workload put on it by the NNM system, and is delivering a great benefit to the company by performing its duties well.

There is room for improvement when it comes to the tracking of NNM web map access by users, and controlling it by authorized and unauthorized users. NNM system patches and NNM server OS and IIS patches should be installed as soon as they are released provided company policy allows mission-critical servers to be patched before tested. This will mitigate the greatest risk to an NNM system, Windows server, and IIS web server. It won't take much effort or time for the NNM administrator to have the NNM system and server as secure as possible for the enterprise after following the recommendations made after each audit test on how to correct the problems found.

Most of the findings from each audit step can be easily corrected because of the step-by-step recommendations that have been supplied. Some audit steps have results that may not be able to be fixed because of a tight budget or enterprise-specific requirements or limitations that keep the NNM administrator from following through with the recommendations that were given. The NNM system is doing well, and will keep performing securely if the recommendations that were made are followed through and continue to be practiced in the future.

4.2 - Audit Findings

Step 3: *Public* should be removed as the SNMP RO string from all devices

Audit Findings: The NNM system was configured correctly to not use the default SNMP RO string of *Public* as the monitored network's globally accepted RO string. Although, *Public* was used as an accepted RO string on a node-by-node basis; which was done to catch misconfigured devices on the network.

Background/Risk: It is risky to allow system administrator's to use *Public* as the default RO string, because *Public* is the first string that an attacker will use when trying to connect to the device using an SNMP utility. If an attacker were able to

56

guess an SNMP RO string, they would be able to access valuable network information, the network topology, or local system configurations of the system they have accessed. SNMP RO and RW strings should be treated as passwords, and should be difficult to guess by an attacker.

Audit Recommendations: It is recommended to set the NNM system to locate SNMP-enabled devices that have *Public* set as their RO string, as it is right now. Once these devices have been located by the NNM system, they should be reported to either the system administrator or the manager of the personnel team that is responsible for that device and the SNMP settings should be changed immediately. The corporate SNMP policy should also directly state that any SNMP-enabled devices that are placed on the network should conform to the defined SNMP configuration process of the enterprise. The policy will cut down on the time spent by the NNM administrator checking for newly found devices on the NNM maps that have *Public* set as the RO string, and having to report that to each system administrator.

Costs: A corporately defined SNMP policy that addresses all issues of a devices SNMP configuration will actually help save money for the company. Instead of the NNM administrator having to confront each system administrator that is responsible for the devices that are found by the NNM system with *Public* set as the RO string, time can be spent on the development of a corporate SNMP configuration policy. Once all system administrators using SNMP have recognized the policy it will begin to self-regulate the SNMP configuration process. The NNM system will still need to be allowed to find any devices with *Public* set as the RO string to keep system administrators accountable to the policy. A corporate SNMP policy with consequences for deviating from the outlined standard procedures will save a lot of time the NNM administrator can use administering the NNM system.

Compensating Controls: It won't take a great deal of time to complete a corporate SNMP policy. It would be similar to the corporate domain password policy. The SNMP RO string should be different for each group of functionally related devices on the network, changed periodically, difficult to guess, difficult to brute-force crack, and kept a secret between personnel groups that aren't involved with the devices. If the drafting and implementation of a new corporate SNMP policy is not an option at this time, the NNM administrator should continue diligently watching for any device to be found by the NNM system that is using *Public* as the RO string.

Step 4: SNMP RW strings should be removed from all SNMP-enabled devices

Audit Findings: I found that the NNM system was setup correctly as a network device monitoring system, which would not be using any SNMP RW strings to manage SNMP-enabled devices on the network. The devices listed under the specific nodes section on the NNM system were not using *Private* as the RW string. There were no nodes found using any SNMP RW string.

Background/Risk: If an SNMP-enabled device on the network had an SNMP RW string configured, an attacker would first try to control the devices using *Private* as the RW string. If the device were not using *Private* as the RW string, they would proceed to guess the RW string or use a brute-force SNMP community string utility to try to discover the RW string automatically over time. Once the RW string is discovered for a device the attacker could then use a utility to control the device using SNMP

57

commands that are preconfigured into an SNMP-enabled device. Once the attacker has found the RW string that is being used on that machine, he can then try to gain control of every other machine that can be found using that same RW string. If there is a trust relationship with the machine that is compromised by the attacker and another machine, that machine may be taken over from the machine that is now controlled by the attacker.

Audit Recommendations: A formal corporate SNMP policy stating the official SNMP configuration of SNMP-enabled devices on the network, should be drafted and put into effect by a network security director. This will be a self-regulating document that states consequences for placing a device with misconfigured SNMP settings on the network. The NNM system should still be reviewed once a week to see if it has detected any device that has been placed on the network that deviates from the desired SNMP settings dictated by the corporate SNMP policy.

Costs: The drafting and adoption of a corporate SNMP policy would actually cut costs for the organization. Without this policy, the NNM administrator must devote time to tracking down and watching the NNM system for misconfigured SNMP devices that are placed on the network with an RW string. This time can be freed up by a self-regulating document that states the desired SNMP configuration for all SNMP-enabled devices. The newly recovered time for the NNM administrator can be devoted to other important configuration and maintenance duties that come with a mission-critical NNM system.

Compensating Controls: The drafting and implementation of an additional or entirely new SNMP corporate policy won't be difficult. If a new SNMP corporate policy was put into place from the first step/recommendation covering RO strings, then an extension to that covering RW strings can be easily added. If an entirely new policy is created from this point, it should be comprised of the topics that were discussed above in the first audit step/recommendation, as well as specific considerations that must be given to RW strings. This area of the document should discuss the need for any SNMP-enabled device to have the RW string removed from its SNMP configuration settings. If an RW string is absolutely necessary for the system administrator of that device to be able to manage it correctly, then the string should conform to the RO string policy steps for hardening and securing an SNMP community string outlined in the RO string recommendations discussed earlier. Permission should also be given for this device to use an RW string, and should be reviewed from periodically to make sure the system administrator is following the documented procedure for using RW strings.

Step 8: Check that NNM root directory Share specifies authorized NNM users

Audit Findings: This objective was successfully completed by the NNM administrator. The root directory of the NNM installation can only be accessed by authorized NNM users via Windows Explorer or an NNM remote management client. NNM remote management clients call for the root directory of the NNM installation path to be Shared to allow authorized users to connect to the NNM maps. By default, when a Windows Share is turned on to allow remote file browsing connections using Windows Explorer, every domain user is allowed to connect with full control unless explicitly defined file-access privilege levels are specified for each authorized user.

58

This precaution had been performed by the NNM administrator.

Background/Risk: This step is critical to securing the NNM system, because if the root directory of the NNM installation is Shared and left with default configurations to allow anyone to connect to it, then the integrity of the NNM database and topology information is in jeopardy. Any user can browse this directory remotely, and can also connect to the NNM system using an NNM remote management client, called an NNM management console. If an attacker were able to connect using either means, the database could be destroyed or stolen, and the attacker would then have intimate knowledge of potentially every SNMP-enabled device and the entire topology of the network, using an NNM console connection.

Audit Recommendations: The current Windows Share configuration of the NNM root installation directory is correct and should be kept in its current state. Only authorized NNM users should have access to this directory Share on the NNM server. File-access privileges higher than read-only should be given to certain users only if absolutely necessary. This list should also be reviewed periodically to validate constant compliance.

Costs: There are no additional costs involved for the compliance of this auditing step. The NNM administrator has correctly configured this objective, and should only review the configuration from periodically.

Compensating Controls: There is nothing that can be added to this process to make it any better or more cost-efficient from its current state. The steps that were performed by the NNM administrator must be done to ensure the security of the NNM server.

Step 9: Lock down web access to specific authorized NNM users

Audit Findings: I was able to open the NNM web map and SNMP tools from a machine on the network by opening a web browser window and then pointing my browser to this address http://$NNMserver/OvCgi/ovlaunch.exe. The NNM web map and SNMP tools came up without prompting me for an authorized username and password.

Background/Risk: The NNM web map and SNMP tools give a user access to view the entire topology of the monitored network, and use advanced SNMP tools to gain system information about particular SNMP monitored devices. An attacker would gain a wealth of knowledge about the topology of the network and specific system information from devices if they could access this site hosted by IIS on the NNM server.

Audit Recommendations: There are specific files that can be setup by the NNM administrator to limit access to the web maps and SNMP tools, which are accessible from any web browser on the local network. The NNM administrator can simply run a script that resides locally on the NNM server, and input authorized usernames and their passwords. The username and password pairs will be encoded in a file that resides in the web area of the NNM configuration files.

Costs: This step in securing the NNM system costs a small amount of the NNM administrator's time. Each authorized NNM user will need to be present while the NNM administrator inputs the username and password pair of each person, but this will take less than a minute for each user.

59

Compensating Controls: There are no alternative actions that can be used to give the NNM system this type of protection other than this hardening step.

Step 15: Verify that management station is running latest major NNM release

Audit Findings: The NNM management station was checked for the version of NNM that it is running. The NNM management station is running NNM version 6.4, which is a stable version, but is not the most current.

Background/Risk: There are two points of view to consider when deciding, which NNM release version should be run in a live monitoring environment. Some NNM administrator's will say that the latest NNM release should not be installed until the first patch has been released for the software, so that any new bugs found will have been repaired before it is installed as a mission-critical system. Others will say that only the latest and greatest software versions are going to be the least vulnerable, because new software will have corrected the vulnerabilities from the previous release. The focus is on having the most stable and secure version of the NNM software released from HP.

Audit Recommendations: I recommend that if the company has a software subscription for new NNM version releases from HP, then the NNM administrator should setup the new software on a test server. A test server can be built and configured quickly, and the NNM software can be setup within a few minutes. Once the NNM administrator has both the current NNM system and the new test system running in parallel and under the same load for a week or so, an informed decision can be made as to whether the new NNM software should be put into production or not. More likely than not, the new NNM software release will run very well in your environment, and the NNM administrator won't hesitate to cutover after a trial period.

Costs: This will be the most expensive security measure to implement. A server will be needed temporarily with a hardware configuration comparable to the current NNM server for comparison of the two NNM systems. A day or two of the NNM administrator's time will need to be dedicated to this software research project.

Compensating Controls: The alternative is to keep the NNM system running on the same version of software until a second full version has been released and then upgrade up one release level. Always running one version behind the latest release is acceptable and has no risk when upgrading. The NNM administrator will need to diligently patch the current software version when vulnerability patches are released.

Step 16: Verify NNM backups run locally, weekly, and are stored remotely

Audit Findings: The NNM system was backed up on a weekly basis, and then the entire NNM server's primary hard drive was backed up on the enterprise's backup solution.

Background/Risk: If this were not being performed on a regular basis, then the NNM system would be at risk of not having a known good-state backup of the NNM configurations in the case of an NNM server hardware failure, or a data compromise from an attack. Full backups of the NNM server's hard drive that are restored after a system crash would not allow the NNM administrator to bring the NNM system back up to a known good state, because of NNM software limitations.

Audit Recommendations: The NNM system should always be checked to make

sure that a weekly backup of the NNM system configurations ran successfully. The enterprise backup solution's configurations should be checked to make sure that they always run directly after the NNM system backups. The hard drive capacity should also be checked to verify that there are multiple GB's of free space left on the hard drive, for OS stability after NNM backups have ran.

Costs: There is no cost to consider for this audit step, other than the time it currently takes the NNM administrator to verify the correct backup settings.

Compensating Controls: There are no other alternatives for this audit step.

Step 17: Check NNM patch level for compliance with most recent release

Audit Findings: The NNM system did not have the most recent patch installed. When the system was checked for compliance to this test, a patch had been released earlier that morning, but because the NNM system patches are always scheduled to be installed during a change control period the NNM system did not pass the test.

Background/Risk: The NNM system should always have the most recently released patches applied to it as soon as possible. HP releases these patches to repair a bug in the code that is either opening up a security vulnerability to the NNM system or because an instability in the code has been detected by several other users. There is risk to not have the latest NNM system patches installed and a risk when installing them. An unpatched NNM system could crash from an attacker exploiting a vulnerability, because the code becomes unstable, or from a damaged patch installation executable that corrupts the NNM installation.

Audit Recommendations: The NNM system patches should always be installed as soon as possible, but a full NNM system backup and an NNM server backup should be performed before a patch is installed. The NNM administrator should also always choose to backup the files that are being overwritten during the installation setup process of each patch.

Costs: There are no additional costs involved, because the NNM administrator is already practicing this audit objective.

Compensating Controls: There are no alternatives to this NNM system hardening process. Diligence in patching the NNM system is the only thing that will accomplish this objective.

Step 21: Check for commonly overlooked NNM performance issues

Audit Findings: There are several related tests that are checked during this auditing step. The NNM system passed all of the tests, except the NNM system did seem to stall on the border routers' scheduled SNMP configuration poll.

Background/Risk: If the NNM system hangs on a particular device's SNMP configuration poll, this can fill up the three available spots in the NNM polling queue. If the NNM system falls behind because of long configuration polls or because the system has too much of a load on its CPU, unreliable status notifications for devices will begin to occur.

Audit Recommendations: Some specific polling configurations should be defined for the border routers that are monitored by the NNM system. Border routers have large routing tables that the NNM system tries to download, that aren't actually needed for topology updates because the topology information held in a border

61

router's memory is outside of the internal monitored network. Time spent on this type of device polling is unnecessary and jeopardizes the reliability of the information given to NNM users for other devices being monitored.

Costs: It won't take long for the NNM administrator to locate devices that may be stalling during their configuration polls. A few polling configuration changes can be set quickly, which will take care of this risk concerning the stability and responsiveness of the NNM system.

Compensating Controls: The other option to correct stalling during configuration checks, is to unmanage the border router icon on the RW map. Additional configuration changes can be made to the NNM system as to how it treats the specific devices during SNMP configuration checks, but this can be a tedious process of minor changes over time.

Step 24: Check to see who has administrative privileges on the NNM server

Audit Findings: The NNM server was configured correctly concerning OS Administrators. Only domain users that the NNM administrator had authorized to have administrative access-level privileges were setup on the OS as Administrators.

Background/Risk: If a user was granted administrative-level privileges to the NNM server's OS, they could then remotely connect to the NNM server using Terminal Services, Microsoft Computer Management, and Windows' filesystem Shares. This kind of remote control can allow a user full control over all data and system configurations of that server.

Audit Recommendations: These settings should be reviewed periodically to ensure that only authorized NNM users that are supposed to be Administrators to the NNM server are setup that way in the OS.

Costs: This will require no additional cost, because the step is already being accomplished with the current controls in place.

Compensating Controls: There are no other options to accomplish this need, other than what has been described.

Step 27: Check NNM server for latest OS and IIS patch installation

Audit Findings: The NNM server was lacking many important service packs and vulnerability patches concerning Windows and the IIS web server.

Background/Risk: The biggest risk to the NNM server is not having the latest Microsoft patches installed for Windows and IIS. There are much fewer vulnerabilities that will allow an attacker to gain full control of an NNM system through a vulnerability in the NNM code compared to the amount of vulnerabilities in the NNM server's OS or web server. NNM system data can be stolen or destroyed if an attacker gains control of the NNM server through a known exploit that has not been patched on the NNM server's OS or web server.

Audit Recommendations: The NNM administrator should be fully aware of the corporate mission-critical server patching policy. The NNM server should be patched as soon as possible, and checked on a weekly basis for a vulnerability patch that may correct a known exploit contained on the NNM server.

Costs: The NNM administrator will need to devote more time to tracking these patch releases and coordinating a process for checking and updating the NNM server on a

62

regular basis.

Compensating Controls: There are no other alternative paths for completing this objective. Different tools can be used, other than the one that was used to test this audit step, but will require the same amount of time and similar process.

Bibliography

On-line References:

- Securing Hewlett Packard OpenView Network Node Manager on HP-UX 11. (2002). SANS Reading Room. Retrieved September 23, 2003 from, http://www.giac.org/ practical/Rich_Antonick_GCUX.doc
- SNMP Security Pack. (2003). SNMP Research International, Inc. Retrieved September 24, 2003 from, http://www.snmp.com/ products/snmpsecpack.html
- Network Node Manager Advanced Edition 7.0 Overview and Features. (2003). HP OpenView. Retrieved September 26, 2003 from, http://openview.hp.com/ products/nnm/index.html
- Software Patches. (2003). HP OpenView Support. Retrieved September 29, 2003 from, http://support.openview.hp.com/ cpe/patches/nnm/6.4x/win.jsp
- Your Greatest Strength can become your Greatest Weakness: Simple Network Management Protocol Vulnerabilities. (2003). SANS Reading Room. Retrieved October 1, 2003 from, http://www.sans.org/ rr/papers/44/376.pdf
- SNMP Brute Force Attack. (2003). SolarWinds. Retrieved October 2, 2003, from, http://www.solarwinds.net/ Tools/Security/SNMP_Brute_Force/index.htm
- Cisco Aironet Access Point default public user. (2003). Internet Security Systems. Retrieved October 4, 2003 from, http://xforce.iss.net/ xforce/xfdb/6296
- Technical Knowledge Base. (2003). HP OpenView. Retrieved October 7, 2003 from, http://openview.hp.com/ sso/ecare/keyword
- Physical Security: first steps to a secured network. (2003) ActivSupport. Retrieved October 9, 2003 from, http://www.activsupport.com/ network/vpn_security/physical_security.html
- Windows 2000/XP: Keeping a record of Windows user login attempts. (2002). My Brandeis. Retrieved October 10, 2003 from, http://my.brandeis.edu/ bboard/q-and-a-fetch-msg?msg_id=0000Ya)

- How To: Configure the size and Behavior of Event Viewer Logs in Windows 2000. (2003). Microsoft Knowledge Base. Retrived October 12, 2003 from, http://support.microsoft.com/ ?kbid=320121
- Microsoft Windows Update. (2003). Microsoft Windows Update. Retrieved October 15, 2003 from, http://v4.windowsupdate.microsoft.com/ en/default.asp
- Proactive Vulnerability Assessments with Nessus. (2002). SANS Reading Room. Retrieved October 17, 2003 from, http://www.sans.org/ rr/papers/5/78.pdf
- Download. (2003). Nessus. Retrieved October 18, 2003 from, http://www.nessus.org/ download.html
- Demonstration. (2003). Nessus. Retrieved October 20, 2003 from, http://www.nessus.org/ demo/index.html

Print References:

- Hewlett-Packard Company (2002). *OpenView Managing Your Network for NNM*. CO: Air Media, Inc.
- Hewlett-Packard Company (2002). *OpenView Scalability and Distribution for NNM*. CO: Air Media, Inc.

Appendix A: Auditing Tools

- SNMP Brute Force Attack
    - SolarWinds
        - http://www.solarwinds.net/Tools/Security/SNMP_Brute_Force/index.htm
- Etherpeek
    - WildPackets
        - http://www.wildpackets.com/products/etherpeek
- Nessus
    - Nessus
        - http://www.nessus.org