>BUSINESS MADE **SIMPLE**

NØRTEL

Nortel Security Solutions
Product Readme

> Nortel Switched Firewall (NSF)
4.1.x - Readme

Document Date: 29 July, 2006
Document Version: 1.4

Table of Contents

Change Log

Version	What	When	Who
1.0	Readme for 4.1.1 release	07/24/2005	Rajesh Vijayakumar
1.1	Added CR Q01113493	08/23/2005	Rajesh Vijayakumar
1.2	Readme for 4.1.2 release	03/22/2006	Ganesh Lakshmanan
1.3	Readme for 4.1.2.1 R61 release	07/07/2006	Ganesh Lakshmanan
1.4	Layout, documentation corrected	07/29/2006	Jürgen Luksch

1 Introduction

This is the consolidated readme for all NSF 4.1.x releases. The objective of a single readme is to help the reader find and track the status and history of an issue more easily. In order to meet this objective, the document is organized in different sections as follows.

Section 2 contains a table that lists the status of all known issues found in 4.1.x releases. It shows the release where the issue was found, the current status of the issue, and the status of the issue in each 4.1.x software release. The following section (Section 3) describes the procedure to upgrade from an earlier release to a 4.1.x release.

The following sections present the detailed readme for each release (one section for each release). These sections describe the hardware platforms and Check Point software versions supported by each release. Finally, the list of all known issues with a brief description and work around (if any) is presented in Appendix A. The current status of each issue is also presented as part of the description.

2 Status of Known Issues and Limitations

All the known issues found and/or fixed in 4.1.x releases are summarized in the following table. The details of the issues are described in Appendix-A. Each row in the table corresponds to a known issue. A detailed explanation of the issue can be found by looking at the CR# (if available) in the Appendix. If a CR# is not available for an item, then search for the issue title in the Appendix for the specific update date. The known issues without the CR# are listed at the beginning of each sub-section for the specific update date.
If you are viewing this document on your computer, you can click on a description item to jump to the full description in Appendix A.

The current status and the status of the issue in different releases are also presented in the table. In the table, ☒ means the particular build is affected and ☑ means the issue is fixed in the particular build.

Table 1   Current status of all issues found in NSF-4.1.x releases.

CR #	Description of Issues and Limitations	Last Updated	Current Status	Status in Different Releases		
				4.1.1	4.1.2	4.1.2.1
Q01338725	HFA is not getting upgraded during software upgrade from 4.1.1_R55 to 4.1.2_R55	03/22/06	Open	☒	☒	☒
Q01338744	mond log file is not getting rotated when we run the system more than 10 days	03/22/06	Open	☒	☒	☒
Q01218241	Primary port failing to copper backup flips cluster	03/22/06	Closed	☒	☑	
Q01229738	snmpd process keeps restarting	03/22/06	Closed	☒	☑	
Q01208951	NTP server access is not restricted to SFD subnet	03/22/06	Closed	☒	☑	
Q01234723	NSF 4.1.1 Replacing a SFA shows 3 SFA's in /info/det	03/22/06	Closed	☒	☑	
Q01248760	unnecessary/confusing log message related to SFD status in single SFD setup	03/22/06	Closed	☒	☑	
Q01279395	Accelerator is 1 hour ahead of directors	03/22/06	Open	☒	☒	☒
Q01326194	L2 Vlan does not work properly	03/22/06	Open		☒	☒
Q01252655	/cfg/sys not available to change NAAP VLAN id on accelerator	03/22/06	Closed	☒	☑	
Q01284910	VPN Site to Site doesn't work	03/22/06	Open		☒	☒
Q01266907	SSI Restarting	03/22/06	Open	☒	☒	☒
Q01296879	bogus errors on ports 27 and 28 of the accelerator	03/22/06	Open	☒	☒	☒
Q01335693	NSF /var/tmp/local_state file filling up	03/22/06	Closed	☒	☑	
Q01113493	FTP session fails when TCP sequence verification is enabled	03/22/06	Open	☒	☒	☒
Q01061470	Sync thro vnic is not supported with VPN	03/22/06	No Fix Planned	☒	☒	☒
Q01050555	Using copper GBIC on dual ports of 6600	03/22/06	No Fix Planned	☒	☒	☒
Q01104775	Proxy Arp entries above 1600 does not work	03/22/06	Open	☒	☒	☒
Q01052781	"Softdog driver open failure" message when accelerator boots up	03/22/06	Open	☒	☒	☒
Q01094577	Pimd restarts when ip address on pim enabled interface is changed	03/22/06	Open	☒	☒	☒
Q01089358	Changes made to routemaps do not take effect immediately	03/22/06	Open	☒	☒	☒

CR #	Description of Issues and Limitations	Last Updated	Current Status	Status in Different Releases		
				4.1.1	4.1.2	4.1.2.1
Q01142037	Issues related to large configurations	03/22/06	Open	☒	☒	☒
Q01117033	Incorrect UDP Blast behavior	03/22/06	Open	☒	☒	☒
Q01157944	Director runs out of memory during bootup	03/22/06	Open	☒	☑	
Q01174716	Dos attack traffic causes high MP CPU utilization on accelerator	03/22/06	Closed	☒	☑	
Q01164785	Error message for pmatch string longer than 40 characters is not clear	03/22/06	Closed	☒	☑	
Q01138787	Limit on number of virtual NICs	03/22/06	Open	☒	☒	☒
Q01159781	Changing subnet on RIP enabled interface requires restarting RIP	03/22/06	Open	☒	☒	☒
Q01158579	Static routes are lost after importing configuration	03/22/06	Open	☒	☒	☒
Q01149215	No validation check to prevent invalid subnet mask for IP ACL	03/22/06	Closed	☒	☑	
Q01133343	/info/sensors show CPU and board temperature to be negative	03/22/06	Open	☒	☒	☒
Q01150870	PIM: NSF does not support fragmented PIM messages	03/22/06	Open	☒	☒	☒
Q01157101	Validation error about accesslist after upgrade	03/22/06	No Fix Planned	☒	☒	☒
Q01160601	"asfcapture" does not capture packets simultaneously on accelerator and director	03/22/06	Closed	☒	☑	
Q01157140	Unable to configure accelerator after upgrade from 4.0.x	03/22/06	No Fix Planned	☒	☒	☒

3 Upgrading to NSF 4.1.x

Upgrade to 4.1.x is supported from 4.0.1 or later versions. 4.1.2.1 R61 requires 250 MBytes free space on the /isd partition. To check available free space, login as root, run "`df -H /isd`" and look under the "Avail" column. If you do not have enough free space, you will get an error saying "`Failed to unpack software…`" when you try to download the .pkg file.

If there is not enough free space for upgrade, please export the current configuration using "`/cfg/ptcfg`", do a clean install from CD, and then import the configuration using "`/cfg/gtcfg`". When configuration exported from 4.0.2 or below is imported into 4.1.1, you will lose all configured static routes. Please see Q01158579 on how to recover the static routes.

When upgrading from 4.0.x to 4.1.x, please keep the following things in mind. 4.1.x is a combined L2/L3 firewall. If you have multiple ports in the same VLAN, the default behavior of 4.1.x is to apply the firewall policy to traffic that is bridged between the ports. This is different from the 4.0.x behavior which applied the firewall policy only to routed traffic. If you would like to keep the 4.0.x behavior, please disable L2 firewall processing on these VLANs using the "`/cfg/net/vlan <n>/l2fw`" CLI item after upgrade.

To upgrade, first download the appropriate 4.1.x upgrade package to the cluster. This can be done over the network using "`/boot/software/download`" or from the CR-ROM using "`/boot/software/cdrom`". Run "`/boot/software/cur`" to make sure the new version was downloaded successfully. You can then activate the new version using "`/boot/software/activate`".

If you are upgrading from 4.0.x, please reset SIC on the directors using "`/cfg/fw/sic`", establish SIC again from SmartDashboard and push policy. If you are upgrading from 4.1.x but to a different version of Check Point, re-establish SIC from SmartDashboard and push policy. If you are upgrading from 4.1.x and to the same version of Check Point you are currently running, the SIC and policy will be preserved.

After upgrade from 4.0.x, please make sure the accelerators are configured by running "`/info/det`". If an error is reported, please see Q01157140 to recover.

The summary of the main steps for upgrading to NSF 4.1.2.1 is given in **Table 2**.

Table 2 Upgrading to NSF 4.1.2.1 R61

From	To	Upgrade Steps
4.0.1-x	4.1.2.1-R61	• Do a clean install using iso image.
4.0.2.0a	4.1.2.1-R61	• Use "`/boot/software/download`" to download R61 upgrade package (NSF_Director_4.1.2.1_R61.pkg). This should be done only in one SFD. • Activate 4.1.2.1 image using "`/boot/software/activate`". This should be done only in one SFD. • Please wait until SFDs reboot and all upgrade process is complete.

		• Get topology in the Check Point management station and push the policy. • Do the post-upgrade verification.
4.0.3.0 4.0.4.0	4.1.2.1-R61	• Use "`/boot/software/download`" to download R61 upgrade package (NSF_Director_4.1.2.1_R61.pkg). This should be done only in one SFD. • Activate 4.1.2.1 image using "`/boot/software/activate`". This should be done only in one SFD. • Please wait until SFDs reboot and all upgrade process is complete. • Get topology in the Check Point management station and push the policy. • Do the post-upgrade verification.
4.1.1	4.1.2.1-R61	• Use "`/boot/software/download`" to download R61 upgrade package (NSF_Director_4.1.2.1_R61.pkg). This should be done only in one SFD. • Activate 4.1.2.1 image using "`/boot/software/activate`". This should be done only in one SFD. • Please wait until SFDs reboot and all upgrade process is complete. • Get topology in the Check Point management station and push the policy. • Do the post-upgrade verification.
4.1.2	4.1.2.1-R61	• Use "`/boot/software/download`" to download R61 upgrade package (NSF_Director_4.1.2.1_R61.pkg). This should be done only in one SFD. • Activate 4.1.2.1 image using "`/boot/software/activate`". This should be done only in one SFD. • Please wait until SFDs reboot and all upgrade process is complete. • Get topology in the Check Point management station and push the policy. • Do the post-upgrade verification.

3.1   Pre-Upgrade Preparation

Backup configuration: You are strongly advised to backup the NSF configuration before doing the upgrade. Please use "`/cfg/ptcfg`" command to export the configuration. This should be done only in one SFD.

3.2   Downloading the upgrade Package

The upgrade package can be downloaded by two different ways. In the first method, the image can be downloaded via FTP using "`/boot/software/download`" CLI command. The CLI will prompt for all the details information, such as IP address of the server and the filename on the server, etc.

Since the NSF installation CD contains the upgrade files (i.e. pkg files), it can be used to import the pkg file to the SFD. User can also burn his/her own CD containing the pkg file. Note that upgrade process requires that file extension to be .pkg. The CD-ROM is automatically ejected at the end of the operation. This step should be done only in one SFD.

3.3   Activating the new software

Once the upgrade package is downloaded, "`/boot/software/cur`" can be used to display all the software versions in the SFD. The version that was just imported will have the status "unpacked." The new version (4.1.2.1) can now be activated using "`/boot/software/activate`". This should be done only in one SFD.

The activation process will upgrade both the Nortel software and the Check Point software to the same version as a clean install from the CD. There is no need to upgrade the Check Point software separately. Each SFD will reboot twice during the upgrade process: once after the upgrade of Nortel software and again after upgrading the Check Point software. The whole process could take somewhere between 15-20 minutes.

After successful software upgrade, the following steps must be done:
- Get topology information in the Check Point management station and
- Push the policy to the NSF cluster.

3.4   Post-Upgrade Verification

The following steps should be done to verify that the upgrade process was completed successfully. These steps are not required for a successful upgrade. However, it is recommended only for the purpose of verification.
- Login as root and run "`os-version`". You will get the output "1.5.1.3_tng.4.1.2.1_R61" for R61.
- Login as admin and check "`/info/cluster`" CLI to make sure that all the directors in the cluster are working fine.

4 Nortel Switched Firewall, 4.1.1 (07/25/2005)

4.1 Supported Hardware Platforms

NSF 4.1.1 supports the following hardware platforms:
- NSF 6414 (6400 accelerator with 5014 director)
- NSF 6614 (6600 accelerator with 5014 director)
- NSF 6424 (6400 accelerator with 5024 director)
- NSF 6624 (6600 accelerator with 5024 director)
- NSF 6416 (6400 accelerator with 5016 director)
- NSF 6616 (6600 accelerator with 5016 director)
- NSF 6426 (6400 accelerator with 5026 director)
- NSF 6626 (6600 accelerator with 5026 director)

4.2 Supported Check Point Releases

NSF 4.1.1 supports the following Check Point versions:
- Check Point NG with Application Intelligence R54 (FP4) with HFA-414.
- Check Point NG with Application Intelligence R55 with HFA-12.

4.3 What's New

The following new features are supported in this release:

4.3.1 Integrated L2/L3 Firewall

NSF 4.1.x can operate either in L2 mode, L3 mode or combined L2/L3 mode depending on the configuration. To configure ports in L2 mode, define a VLAN and add the ports to the VLAN. Also make sure L2 firewall processing is enabled for the VLAN using "`/cfg/net/vlan <n>/l2fw`" CLI or "Config | Network | VLANs" page in WebUI. L2 firewall processing is enabled by default.

To configure the system in L3 mode, please disable L2 firewall processing for all VLANs that have multiple ports in them. For VLANs with only one associated port and for automatic VLANs, L2 firewall is always off.

To configure a VLAN in combined L2/L3 mode, define the VLAN, add at least 2 ports to the VLAN, enable L2 firewall processing for the VLAN, define an interface and add it to that VLAN. All bridged traffic between ports in the VLAN will go through L2 firewall processing and all routed traffic to another VLAN will go through L3 firewall processing.

4.3.2 Security Pack

Security Pack detects and prevents a number of attacks right on the accelerator.

- Protects against common DoS attacks like Smurf, LandAttack, Fraggle, NullScan, XmasScan, ScanSynFin, PortZero and Blat. This can be configured using "`/cfg/net/port <n>/sec/dos`" from the CLI or "Config | Network | Ports | Security" page in WebUI.

- Management processor rate limiting protects the MP on the accelerator by limiting the number of packets forwarded to it. This can be configured using "`/cfg/acc/mprlimit`" menu in the CLI or "Config | Cluster | Accelerator(s) | General" page in the WebUI.

- UDP blast protection protects UDP services from attack by limiting the amount of traffic on a per service basis. UDP blast settings can be configured using "`/cfg/sec/udpblast`" in the CLI or "Config | Security | UDP Blast" in the WebUI. UDP blast protection can be enabled on a per port basis using "`/cfg/net/port <n>/sec/udpblast`" in the CLI or "Config | Network | Ports | Security" page in WebUI.

- IP range access list (IP ACL) allows user to configure up to 5000 IP addresses to be blocked, for example an ISP black list. IP ACL can be configured using "`/cfg/sec/ipacl`" from the CLI or "Config | Security | IP ACL" from the WebUI and can be enabled on a per port basis using "`/cfg/net/port <n>/sec/ipacl`" in the CLI or "System | Network | Ports | Security" page in WebUI.

- Protocol rate limiting allows you to rate limit TCP, UDP or ICMP sessions. When the threshold rate is exceeded, new sessions will be dropped until the configured hold-down period exceeds. To use protocol rate limiting, you first define a filter using "`/cfg/net/adv/filt`" CLI or "Config | Network | Filters | Filters" page in WebUI. Then you can define how traffic matching that filter should be rate limited using "`/cfg/net/adv/filt <n>/adv/rlimit`" menu in CLI or "Config | Network | Filters | Rate Limiting" page in WebUI. Finally, apply the filter to specific ports using "`/cfg/net/port <n>/enf`" and "`/cfg/net/port <n>/filt`" in CLI or "Config | Network | Ports | General | Modify" page in WebUI.

- Pattern matching allows you to define filters that match the incoming packets against a simple string or a regular expression. To use pattern matching, first define a filter using "`/cfg/net/adv/filt`" CLI or "Config | Network | Filters | Filters" page in WebUI. Then you can define what pattern to search for using "`/cfg/net/adv/filt <n>/adv/pmatch`" menu in CLI or "Config | Network | Filters | Pattern Matching" page in WebUI. Finally, apply the filter to specific ports using "`/cfg/net/port <n>/enf`" and "`/cfg/net/port <n>/filt`" in CLI or "Config | Network | Ports | General | Modify" page in WebUI.

4.3.3 Hitless Upgrade

If you have a high availability setup, consisting of 2 accelerators and 2 or more directors, you can upgrade the cluster with virtually no downtime. To start the hitless upgrade process, please use the "`/boot/software/hitless/activate`" CLI. For hitless upgrade to work smoothly, make sure the following conditions are met.

- Both the active and backup accelerators should have all the network links up
- Do not disconnect any network cables or reboot any accelerator or director while hitless upgrade is in progress.

Hitless upgrade works by upgrading one side of the cluster first, then failing over traffic to that side and upgrading the other side. If you are upgrading between different versions of Check Point, hitless upgrade will pause after upgrading one side and wait for you to establish trust and push policy to the upgraded side before failing over to that side.

Stateful session failover is not available during hitless upgrade because Check Point sync will not work between different versions.

4.3.4 Packet Capture

NSF 4.1.x includes a packet capture utility to troubleshoot traffic related issues. It captures the packet at various points within the system as the packet flows through NSF. It supports ethereal like capture filters using the `-f` flag. To start the packet capture utility, login as root and run "`asfcapture -f <filter>`". Please run "`asfcapture -h`" for a brief help message.

4.3.5 Multicast Routing

NSF 4.1.x supports PIM-SM and IGMPv2 multicast routing. NSF can be used as a transit router or Rendezvous Point (RP). In this initial release, NSF cannot be used as an edge router for multicast. This means that multicast receivers and sources cannot be directly attached to NSF. NSF supports up to 31 outgoing VLANs per (S, G) and up to 63 multicast routes.
PIM can be configured using "`/cfg/net/routes/pim`" from the CLI or "Config | Network | Routes | PIM" from the WebUI.

4.3.6 WebUI Enhancements

The WebUI has a new look and feel with a navigation tree pane on the left which makes it easier to move around. The "Wizards" tab has a number of wizards which walk the user through various configuration tasks. There is also an "Initial configuration wizard" that can be launched immediately after initializing a cluster using the "new" item in the CLI, provided you specify a management interface IP address during "new".

The WebUI also includes a "Ticker", which is a Java applet that can be launched from the WebUI. The ticker displays all information needed to monitor NSF on a single screen and can also display charts over a period of time. Java 1.4.2._01 or newer is required for the ticker applet.

4.3.7 Detailed Explanation of Log Messages

Detailed online help is available for various syslog messages generated by the system. Each message contains an identifier (e.g. `CFGD_011`) which can be looked up from the CLI or WebUI to get more details about the message, possible causes and information on how to resolve it. This can be accessed in the CLI using "`/maint/logdetail`" and from the WebUI by navigating to "Config | Administration | Monitor | Syslog" page.

4.3.8 Enhanced Director Load balancing

NSF 4.1.x allows the user to select the load balancing metric to be used for load balancing traffic across the directors.

- iphash – Traditional load balancing metric using the source and destination IP addresses to select the director.
- ipporthash – Use the source port and destination port in addition to source IP and destination IP. Use this if a large portion of the traffic has the same source and destination IP.

You can also specify a weight between 0 and 15 for each director. You can also specify a weight for the MIP holder, which will override the weight specified for that director.

The load balancing options can be configured using "`/cfg/sys/lbopts`" menu in the CLI or "Config | Administration | Load Balancing" page in WebUI..

4.3.9 Gateway Persistency

When using multiple default gateways and round robin metric for load balancing the default gateways, gateway persistency can be enabled to ensure symmetric routing. Gateway persistency can be enabled on a port using "`/cfg/net/port <n>/gwp`" menu item in CLI or "Config | Network | Ports | General | Modify" page in WebUI.

4.3.10 UPS Support

NSF 4.1.x supports UPS devices manufactured by APC Corp. The directors can communicate with the UPS, detect power failure and shutdown cleanly before the UPS battery runs out. NSF can communicate with the UPS over USB cable or SNMP. UPS settings can be configured using "`/cfg/sys/ups`" menu in CLI or "Config | Administration | APC UPS" page in WebUI.

4.3.11 USB Device Support

You can use USB stick storage devices to transfer files from and to NSF. This includes transferring tsdumps, creating system backups, restoring from backups etc.

4.3.12 RADIUS Authentication

NSF 4.1.1 allows you to use an external RADIUS server for authentication for administrative login to NSF CLI or WebUI. RADIUS server authentication can be configured using "`/cfg/sys/adm/auth`" menu in the CLI or "Config | Administration | RADIUS" page in WebUI. The user name must exist on both NSF and the RADIUS server for the user to login successfully. It is recommended that the "fallback" option be enabled so you can login using the local username and password in case the RADIUS server is down.

4.3.13 OSPF Route Maps

Route maps allow you to exercise fine grained control over route redistribution in OSPF. You can define route maps using "`/cfg/net/route/rmap`" in the CLI or "Config | Network | Routes | RMAP" page in WebUI. These route maps can then be attached to "static", "connected" and "RIP" redistribution from OSPF.

4.3.14 Accelerated Sequence Number Verification

Sequence number verification is part of Check Point SmartDefense. Starting with 4.1.1, the accelerator supports sequence number verification. As a result, sequence number verification is done even for accelerated packets. Sequence number verification can be enabled or disabled from the SmartDefense tab of Check Point SmartDashboard.

4.3.15 SCP/SFTP Support

NSF 4.1.1 allows you to transfer files securely over the network using scp or sftp protocol. This can be used to export configuration, tsdumps, backup file etc.

5 Nortel Switched Firewall, 4.1.2 (03/24/2006)

5.1 Supported Hardware Platforms

NSF 4.1.2 supports the following hardware platforms:
- NSF 6414 (6400 accelerator with 5014 director)
- NSF 6614 (6600 accelerator with 5014 director)
- NSF 6424 (6400 accelerator with 5024 director)
- NSF 6624 (6600 accelerator with 5024 director)
- NSF 6416 (6400 accelerator with 5016 director)
- NSF 6616 (6600 accelerator with 5016 director)
- NSF 6426 (6400 accelerator with 5026 director)
- NSF 6626 (6600 accelerator with 5026 director)

5.2 Supported Check Point Releases

NSF 4.1.2 supports the following Check Point versions:
- Check Point NG with Application Intelligence R54 (FP4) with HFA-417
- Check Point NG with Application Intelligence R55 with HFA12 Or HFA-16
- Check Point NGX (R60)

5.3 What's New

There is no new feature added in 4.1.2.

5.4 Configuration of the Gateway Cluster Object for R60

Please refer to Check Point user guide for a detailed description of the procedure to configure R60 SmartDashboard. The following guidelines should be followed while configuring SmartDashboard for NSF:
- While creating cluster object, both VPN as well as ClusterXL in the "Gateway Cluster Properties" window are selected by default. Make sure to *unselect* ClusterXL from the list of Check Point products. Also, *unselect* VPN if it is not used.
- While defining the gateway cluster for the NSF in Check Point SmartDashboard, the "3rd Party Configuration" in the gateway cluster properties should be configured as follows:
  - Cluster Operation Mode: **Load Sharing** (mandatory)
  - 3rd Party Solution: **OPSEC** (mandatory)
  - Support non-sticky connections: **Yes** (mandatory)
  - Hide Cluster Members' outgoing traffic behind Cluster's IP Address: **No**
  - Forward Cluster's incoming traffic to Cluster Members' IP Address: **No**
  - Configure the Check Point synchronization interface in the topology page. This configuration used to be under "Synchronization" tab in "Gateway Cluster Properties" window for R54 and R55.

5.5 Bugs Fixed Since 4.1.1 Release

6 Nortel Switched Firewall, 4.1.2.1 (07/07/2006)

6.1 Supported Hardware Platforms
NSF 4.1.2.1 supports the following hardware platforms:
- NSF 6414 (6400 accelerator with 5014 director)
- NSF 6614 (6600 accelerator with 5014 director)
- NSF 6424 (6400 accelerator with 5024 director)
- NSF 6624 (6600 accelerator with 5024 director)
- NSF 6416 (6400 accelerator with 5016 director)
- NSF 6616 (6600 accelerator with 5016 director)
- NSF 6426 (6400 accelerator with 5026 director)
- NSF 6626 (6600 accelerator with 5026 director)

6.2 Supported Check Point Releases
NSF 4.1.2.1 supports the following Check Point versions:
- Check Point NGX (R61)

6.3 What's New
There is no new feature added in 4.1.2.1

6.4 Bugs Fixed Since 4.1.2 Release

Appendix A: List of Known Issues

This Appendix provides detailed explanation on all the issues found and/or fixed in 4.1.x releases. The following information is provided for each issue:
- Last update date
- Affected releases
- Current status
- Description of the problem
- Description of the work around or fix, if available

Issues Updated on 03/22/2006

HFA IS NOT GETTING UPGRADED DURING SOFTWARE UPGRADE FROM 4.1.1_R55 TO 4.1.2_R55

CR# Q01338725
Last Updated: 03/22/2006
Affected Releases: 4.1.x
Current Status: Open

HFA from 12 to 16 is not getting upgraded during software upgrade from
4.1.1.0a_R55 to 4.1.2.0c_R55. The clean installation of 4.1.2.0c_R55 from CD
has HFA 16. The upgrade from 4.0.x to 4.1.x does not have any issues with HFA upgrade and it comes up with HFA16.

The workaround is to download HFA16 or latest Nortel certified Check Point HFA for NSF platform from Nortel web site and install on each SFD.

MOND LOG FILE IS NOT GETTING ROTATED WHEN WE RUN THE SYSTEM MORE THAN 10 DAYS.

CR# Q01338744
Last Updated: 03/22/2006
Affected Releases: 4.1.x
Current Status: Open

The service mond writes logs to `mond.log` in `/var/tmp` location and when the file size reaches 100MB, the file is not rotated and keeps growing till the service is restarted or SFD is rebooted.
The workaround is to restart mond service periodically which automatically rotates the files at the service startup.

PRIMARY PORT FAILING TO COPPER BACKUP FLIPS CLUSTER.

CR# Q01218241
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: Closed

When the primary physical link of a port failed (fibre), the backup physical link (copper) was activated but a VRRP failover on the accelerator was triggered. By allowing enough time for port failure, vrrp failover was avoided.

SNMPD PROCESS KEEPS RESTARTING.

CR# Q01229738
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: Closed

When the SSI process went down, SNMP went down and health check restarted snmpd process. The fix made the snmp daemon independent of SSI and does not go down when SSI internally restarts. The snmp daemon service is restarted once in every week.

NTP SERVER ACCESS IS NOT RESTRICTED TO SFD SUBNET.

CR# Q01208951
Last Updated: 03/22/2006
Affected Releases: 4.1.2
Current Status: Closed

NTP server access was not restricted to SFD subnet. Technically, NTP server runs on MIP and other SFDs and accelerators access NTP server. Since iptables rule was missing, any host with access to MIP could talk to NTP and get details such as OS info, time etc. The fix was to make NTP accessible within SFD subnets and configured NTP servers.

NSF 4.1.1 REPLACING A SFA SHOWS 3 SFA'S IN /INFO/DET.

CR# Q01234723
Last Updated: 03/22/2006
Affected Releases: 4.1.2
Current Status: Closed

After replacing an accelerator, 3 SFA'a are displayed under CLI command `/info/det`. The fix was to display only detected/reachable accelerators.

UNNECESSARY/CONFUSING LOG MESSAGE RELATED TO SFD STATUS IN SINGLE SFD SETUP

CR# Q01248760
Last Updated: 03/22/2006
Affected Releases: 4.1.2
Current Status: Closed

When there is only one SFD-Accelerator (standalone) in the setup, the process which was invoked as part of determining isolation case from the cluster reported misleading information which was not appropriate for standalone setup. The fix was to check for type of cluster before writing log messages.

ACCELERATOR IS 1 HOUR AHEAD OF DIRECTORS

CR# Q01279395
Last Updated: 03/22/2006
Affected Releases: 4.1.x
Current Status: Open

When the SFD date changes to Apr 02 2006 01:59:00 (timezone: North America), the SFD time changes to 3AM (PST) after one minute but the accelerator still shows 2AM PST.

L2 VLAN DOES NOT WORK PROPERLY

CR# Q01326194
Last Updated: 03/22/2006
Affected Releases: 4.1.2 R60, 4.1.2.1 R61
Current Status: Open

When a vlan is configured with two ports and L2FW is enabled for the vlan, CheckPoint NGX drops traffic in the vlan. This issue happens only in 4.1.2 R60 and 4.1.2.1 R61.

/CFG/SYS NOT AVAILABLE TO CHANGE NAAP VLAN ID ON ACCELERATOR

CR# Q01252655
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: Closed

User is unable to set NAAP VLAN id in accelerator on NSF 4.1.1 since the `/cfg/sys` menu is not available in the accelerator when logged in as 'admin' user. The issue is fixed and now user can change default NAAP vlan. The `/cfg/vlan` menu is hidden in 'admin' mode since user should not accidentally change vlans in the accelerator.

VPN SITE TO SITE DOESN'T WORK

CR# Q01284910
Last Updated: 03/22/2006
Affected Releases: 4.1.2 R60
Current Status: Open

Both site-to-site and client-to-site VPN will not work in 4.0.4-R60 with the default management station settings. The work around to resolve this problem is described below:

1. Configure the VPN gateway object in the Check Point SmartDashboard and save the configuration. Close the management station if it is opened.
2. Open a dos window.
3. Type "`cd \program files\checkpoint\smartconsole\r60\program`". If you have installed Check Point management software in a different location, you should cd to appropriate directory.
4. Type in "`guidbedit`" and connect to management station.
5. Hit "`Ctrl F`" (for find) and type "`reroute`" in the "Find What" box
6. Click on the "Find Next" button
7. It should take you to the "reroute_encryted_packets" in the "Field Name" column
8. Change the "Value" to false.
9. Hit "F3" and it should find the next instance of "reroute_encryted_packets"
10. Change its "Value" to false.
11. click on "File" and click on "Save All"
12. Close the guidbedit window and start it again and double check the values of "reroute_encryted_packets" are set to false.
13. Close the guidbedit window after verifying the values.
14. Start the management station and push the policy to the firewall.

In addition, if the encryption domain is NAT'ed and VPN community is used, it may be necessary to disable NAT inside the VPN community. The Disable NAT inside the VPN Community property checkbox can be toggled in the SmartDashboard (VPN Manager tab -> Community object properties -> Advanced VPN Properties tab). Disabling the reroute_encrypted_packets property for a NPV community also prevents Excluded Services within the VPN from working. The Excluded Services tab is also inside SmartDashboard (VPN Manager tab -> Community object properties).

SSI RESTARTING

CR# Q01266907
Last Updated: 03/22/2006
Affected Releases: 4.1.x
Current Status: Open

The SSI service restarts intermittently and causes daemons such as clicbd, cfgd, hc to disconnect from SSI and reconnect again. The impact is when user is logged-in to CLI/BBI and restart takes place, the user CLI/BBI is terminated and user needs to connect again.

BOGUS ERRORS ON PORTS 27 AND 28 OF THE ACCELERATOR

CR# Q01296879
Last Updated: 03/22/2006
Affected Releases: 4.1.x
Current Status: Open

The error counters go up on ports 27 and 28 even when nothing is plugged into the ports. The error s only appears at the accelerator ports 27 & 28 and they appear only when both SFA and SFD reboot at the same time.

`/stats/port 27/clear` would clear this statistics.

NSF /VAR/TMP/LOCAL_STATE FILE FILLING UP

CR# Q01335693
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: Closed

When cfgd started, it created dynamic registry nodes in `/var/tmp/local_state` and the last node which was written disappeared. When cfgd tried to access it again, it caught an exception and exited. Health check restarted cfgd which triggered ACCEL OFF and ACCEL ON states on the accelerator.

The issue was the local_state file quickly filled up and the file was rotated. The rotation of the new file was the reason for exception and it was fixed.

Issues Updated on 08/23/2005

FTP SESSION FAILS WHEN TCP SEQUENCE VERIFICATION IS ENABLED

CR# Q01113493
Last Updated: 08/23/2005
Affected Releases: 4.1.1
Current Status: Open

When TCP sequence verification is enabled in Check Point SmartDefense, the FIN packet of the FTP data session is incorrectly dropped with a Check Point log message indicating TCP sequence check failure. With some FTP clients, this will cause the FTP session to hang while the client waits indefinitely for the data connection to close.
The workaround is to turn off TCP sequence verification. Open SmartDashboard and from the SmartDefense tab, go to "Network Security | TCP" and uncheck "Sequence Verifier". Save the changes and push policy to NSF again.

Issues Updated on 07/25/2005

SYNC THRO VNIC IS NOT SUPPORTED WITH VPN
CR# Q01061470
Last Updated: 07/25/2005
Affected Releases: 4.1.x
Current Status: No Fix Planned

If VPN is selected in the cluster object properties of the NSF cluster object in Check Point SmartDashboard, you should not use "sync thro vnic" as this will cause VPN traffic to fail occasionally. Please use a dedicated port on the director for Check Point sync.

USING COPPER GBIC ON DUAL PORTS OF 6600
CR# Q01050555
Last Updated: 07/25/2005
Affected Releases: 4.1.x
Current Status: No Fix Planned

The 6600 accelerator has dual connecters for ports 3, 4, 5 and 6. You can use either copper of fiber GBICs for the GBIC slots. If you use copper GBICs, link negotiation happens between the accelerator and the GBIC causing the accelerator to consider that link as up as soon as the copper GBIC is inserted. If GBIC is configured as you preferred connector, the accelerator will switch over to the copper GBIC as soon as it is inserted even of there is no cable connected. To prevent this, you should set the GBIC port as the backup if you plan to use copper GBICs.

PROXY ARP ENTRIES ABOVE 1600 DOES NOT WORK

CR# Q01104775
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

If you have more than 1600 proxy arp addresses defined, the system will not do proxy arp for entries above the first 1600.

"SOFTDOG DRIVER OPEN FAILURE" MESSAGE WHEN ACCELERATOR BOOTS UP

CR# Q01052781
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

When the accelerator is booting up, you will see the above error message if you are connected to the serial console of the accelerator. The message is harmless and can be safely ignored.

PIMD RESTARTS WHEN IP ADDRESS ON PIM ENABLED INTERFACE IS CHANGED

CR# Q01094577
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

If the IP address of a PIM enabled interface is changed, it will cause the pimd process to restart. This may cause disruption of multicast traffic for a few seconds.

CHANGES MADE TO ROUTEMAPS DO NOT TAKE EFFECT IMMEDIATELY

CR# Q01089358
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

If more than one redistribution is enabled, changes made to routemaps do not take effect even after the changes are applied. As a workaround, please disable and enable OSPF redistributions when routemaps are updated.

ISSUES RELATED TO LARGE CONFIGURATIONS

CR# Q01142037
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

If the NSF configuration is very large, the system will take a long time to apply the configuration or may fail to apply the configuration. Other symptoms include "accel off", failing CLI commands under "/info" menu and high CPU usage on the director as it tries to apply the configuration. Examples of large configurations include 1500+ proxy arp addresses, 4000+ static routes and 2000+ filters.

INCORRECT UDP BLAST BEHAVIOR

CR# Q01117033
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

The CLI allows user to specify multiple UDP ports and port ranges to be protected against UDP blast. However the system applies the UDP blast control to all ports that are between the lowest and highest user specified ports. For example, if you configure UDP blast protection for ports 1000 and 2000, all UDP ports between 1000 and 2000 also get UDP blast protection.
Under heavy stress, when traffic is more than 50% of gig line rate, the system is unable to enforce the configured UDP blast protection (Q01157980).

DIRECTOR RUNS OUT OF MEMORY DURING BOOTUP

CR# Q01157944
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: Closed

If Check Point sync is enabled, the director tries to sync it's session table with that of the cluster members during bootup. However, if session rate is more than 4000 per second and all the sessions are being synched, the director will run out of memory trying to sync up its session table.
FW flags were added to fwkern.conf which limits memory usage for sync at startup.

DOS ATTACK TRAFFIC CAUSES HIGH MP CPU UTILIZATION ON ACCELERATOR

CR# Q01174716
Last Updated: 02/22/2006
Affected Releases: 4.1.1
Current Status: Closed

When DOS attack protection is enabled on a port, the accelerator will generate one syslog message per 128 attack packets dropped. However, if the number of attack packets is large enough, the amount of syslogs generated is enough to overwhelm the MP. Release 4.1.2 has CLI command which sets limit on logs generated per second.

ERROR MESSAGE FOR PMATCH STRING LONGER THAN 40 CHARACTERS IS NOT CLEAR

CR# Q01164785
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Closed

When configuring pattern matching string under "`/cfg/net/adv/filt <n>/adv/pmatch`" menu, there is a limit of 40 characters on the length of the pmatch string. If the length of the configured pmatch string is more than 40 characters, you will get the following error when you try to apply the configuration:

```
Update failed: Match String: Unable to set registry value: bad type of argument in
set operation
```

The validation checks for maximum string length and fixed in 4.1.2.

LIMIT ON NUMBER OF VIRTUAL NICS

CR# Q01138787
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

NSF has a limit of 252 on the number of virtual NICs supported. Each L3 VLAN translates to one VNIC. Each port in an L2 VLAN translates to a VNIC. If your configuration results in more than 252 VNICs, all traffic coming to VNICs above 252 will be dropped.

CHANGING SUBNET ON RIP ENABLED INTERFACE REQUIRES RESTARTING RIP

CR# Q01159781
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

If the subnet mask on a RIP enabled interface is changed, the RIP daemon will continue advertising the old subnet mask. Please login as root on the MIP director and run "`service ripd restart`" to restart the RIP daemon and force it to stop advertising the old subnet mask.

STATIC ROUTES ARE LOST AFTER IMPORTING CONFIGURATION

CR# Q01158579
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

If you exported a configuration from 4.0.2 or earlier using "`/cfg/ptcfg`" and later imported it into 4.1.1 using "`/cfg/gtcfg`", the static routes will not be restored. This is because of the difference in the static route structure between the versions. To recover the static routes, login as root on the MIP director and run "`/opt/tng/bin/post-upgrade --routes –verbose`".

NO VALIDATION CHECK TO PREVENT INVALID SUBNET MASK FOR IP ACL

CR# Q01149215
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

While configuring IP ACL, 0.0.0.0 is not considered a valid subnet mask. However, there is no validation to prevent user from configuring 0.0.0.0 as a subnet mask for IP ACL entry. This invalid configuration will be rejected by the accelerator and you will see the following error under "`/info/det`".

```
IP subnet mask cannot be set to 0.0.0.0
```

To fix the problem, please change the subnet mask to a valid one and apply the configuration.

/INFO/SENSORS SHOW CPU AND BOARD TEMPERATURE TO BE NEGATIVE

CR# Q01133343
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

"/info/sensor" CLI command displays the hardware monitoring information like CPU temperature, motherboard temperature, fan speeds etc. Sometimes, it may not be able to get the correct values and will end up displaying "-1" for CPU and motherboard temperatures. You can recover from this situation by rebooting the director. If the problem recurs, please disable hardware monitoring by logging in as root and running the following commands.

```
make-part-rw / on
chkconfig sensord off
reboot
```

PIM: NSF DOES NOT SUPPORT FRAGMENTED PIM MESSAGES

CR# Q01150870
Last Updated: 07/25/2005
Affected Releases: 4.1.1
Current Status: Open

When using PIM, fragmented "join" or "prune" messages are not supported. Fragmented bootstram message (BSM) is also not supported (Q01150866) . Please ensure that these PIM control messages do not get fragmented.

VALIDATION ERROR ABOUT ACCESSLIST AFTER UPGRADE

CR# Q01157101
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: No Fix Planned

If you upgrade from 4.0.x to 4.1.1, you will get the following validation error when you try to make some configuration change and apply.

```
Invalid setting for /cfg/sys/accesslist.
As an accesslist has been configured the MIP and
all hosts has to be part of the accesslist.
```

This is a new requirement in 4.1.1 that the NSF internal subnet has to be added to the accesslist. This is automatically done for you during clean install. In upgrade case, please add the NSF internal subnet to the accesslist to fix the above error.

"ASFCAPTURE" DOES NOT CAPTURE PACKETS SIMULTANEOUSLY ON ACCELERATOR AND DIRECTOR

CR# Q01160601
Last Updated: 03/22/2006
Affected Releases: 4.1.1
Current Status: Closed

The packet capture utility, "asfcapture", captures packets only on the director even if the command line specifies both the director and accelerator as capture locations. The fix was to modify capture options from "sw" to "all" when accelerators and directors capture is enabled.

UNABLE TO CONFIGURE ACCELERATOR AFTER UPGRADE FROM 4.0.X

CR# Q01157140
Last Updated: 03/22/2006
Affected Releases: 4.1.x
Current Status: No Fix Planned

When upgrading from 4.0.x to 4.1.x, please make sure that the system was able to successfully configure the accelerators after upgrade is complete. You can do this by running "/info/det" and making sure the status of each accelerator says "Accelerator is configured and unicast/igmp/pim routes are uptodate". If the accelerator was not configured, please login to the accelerator CLI as 'admin', set the accelerator to boot with factory default config using "/boot/conf fact" and reboot the accelerator using "/boot/reset". This extra step is necessary because of changes in factory default configuration between 4.0.x and 4.1.x software.