NORTEL SWITCHED FIREWALL 6000 Series
Software Release 4.2.2

1. Release Summary

Release Date : June 2008
Purpose : Software maintenance release to address customer software issues.

2. Important Notes Before Upgrading to This Release

Upgrade to 4.2.2 is supported from 4.1.1 or later versions. 4.2.2 requires 500 MBytes free space on the /isd partition. To check the 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 4.2.2, please keep the following things in mind. 4.2.2 is a combined L2/L3 firewall. If you have multiple ports in the same VLAN, the default behavior of 4.2.2 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. 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.

Procedure to upgrade from CLI

- Use "/boot/software/download" to download R60 or R65 upgrade package NSF_Director_4.2.2.0_R60.pkg or NSF_Director_4.2.2.0_R65.pkg).
- Activate 4.2.2 image using "/boot/software/activate".
- This should be done only in one SFD.
- Please wait until SFDs reboot and all upgrade process is complete.
- Re-establish the trust for each director by,
    - Reset sic on the firewall director (/cfg/fw/sic).
    - Unload the default policy on the firewall director (/maint/diag/uldplcy).
    - On the CP management server, Reset and re-initialize sic on the firewall director object

Notes: *1. Upgrade from 4.0.1-x to 4.2.2 is not supported. Please perform a clean install using .iso image. For all later versions the above procedure works.*
2. Upgrade to 4.2.2 from any previous versions is not supported through BBI. Only CLI upgrade is supported.

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.

Downloading the upgrade Package

The upgrade package can be downloaded in different ways. In the first method, the image can be downloaded via FTP using "/boot/software/download" CLI command. The CLI will prompt all the detailed 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 the file extension to be .pkg. The CD-ROM gets automatically ejected at the end of the operation. This step should be done only in one SFD.

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 (4.2.2) 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. Each SFD will reboot twice (if it is a HA setup) during the upgrade process: once after the upgrade of Nortel software and again for sync to start. The whole process could take somewhere between 15-20 minutes.

After the successful software upgrade, the following steps must be done:
Re-establish the trust for each director by,

a. Reset sic on the firewall director (/cfg/fw/sic).
b. Unload the default policy on the firewall director (/maint/diag/uldplcy).
c. On the CP management server, Reset and re-initialize sic on the firewall director object.

Push the Check Point Firewall policy from the CP management server.

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 4.2.2_R60" or "1.5.1.3_tng.4.2.2_R65"
- Login as admin and check "/info/cluster" CLI to make sure that all the directors in the cluster are working fine.

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. Hitless upgrade will pause after upgrading one side and wait for you to re-establish the trust and push the 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

Procedure for HITLESS upgrade from CLI

- Use "/boot/software/download" to download R60 or R65 upgrade package (NSF_Director_4.2.2.0_R60.pkg or NSF_Director_4.2.2.0_R65.pkg).
- Activate 4.2.2 image using "/boot/software/hitless/activate". This should be done only in one SFD.
- Once upgrade is done to one side of the cluster please perform the following on the firewall director and the CP management server for the firewall to become operational and upgrade to continue to the other side,
    - Reset sic on the firewall director (/cfg/fw/sic).
    - Unload the default policy on the firewall director (/maint/diag/uldplcy).
    - On the CP management server, Reset and re-initialize sic on the firewall director object.
- Push the Check Point Firewall policy from the CP management server.
- Once the other side is upgraded please perform steps 3 & 4 for HA to become operational

Notes: *Upgrade to 4.2.2 from any previous version is not supported through BBI. Only CLI upgrade is supported*

3. Platforms Supported

Hardware Platforms Supported

PEC	MODEL #
EB1639173(E5)	Nortel Switched Firewall system 6416
EB1639174(E5)	Nortel Switched Firewall System 6616
EB1639067(E5) + EB1639131(E5)	Nortel Switched Firewall System 6426
EB1639113(E5) + EB1639131(E5)	Nortel Switched Firewall System 6626
EB1639067(E5)	NSF Accelerator 6400
EB1639113(E5)	NSF Accelerator 6600
EB1639130(E5)	NSF Director 5016
EB1639131(E5)	NSF Director 5026

Supported Check Point applications

Nortel Switched Firewall 6000 Series 4.2.2 supports the following *Check Point applications:
- Firewall-1®
- VPN-1®
- SmartDefense™
- NAT
- Authentication
- Content Security
- Policy Server
- Management Tools
- SmartView Monitor™

Configure the following management tools outside the NSF 4.2.2 software:
- SmartDashboard™
- SmartView Tracker™
- SmartView Status™

4. Notes for Upgrade

File Names for This Release

File Name	Module Or File Type	File Size(K.B)
NSF_Director_4.2.2.0_R60.iso	.iso (contains .img and .pkg files)	346336
NSF_Director_4.2.2.0_R65.iso	.iso (contains .img and .pkg files)	474560

File Name	MD5 Checksum
NSF_Director_4.2.2.0_R60.iso	a332bf61d576ab3f52134242e8b44ff7
NSF_Director_4.2.2.0_R65.iso	ef4adc13fe3dec65cdabbb5794586740

5. Version of Previous Release

Software version 4.2.1, Release date - March 31, 2008

6. Compatibility

N/A

7. Changes in This Release

Problems resolved in this release: 4.2.2

Q01827696 NSF allows users to configure same IP address for two different interfaces (though with different mask.). In iSD, the configuration is getting applied successfully but failed to get applied at the accelerator. Gradually accelerators move to Accel-OFF state and no more configurations can be applied at the accelerator. This issue is fixed in 4.2.2 release by adding a validation not to allow same IP addresses for different interfaces.

Q01834805 After configuring the management IP and sync interface, the directors will go for a reboot. Once the directors come up and if no accelerators are connected at that time, the management IP configuration is lost. This is caused due to a validation in the

configuration module which returns without applying the config if no active accelerators are found. This issue is fixed in 4.2.2 release by removing the validation in the configuration module

Q01855791 L2FW over trunked ports doesn't work as desired and regular traffic is affected. This issue is caused by invalid FDB lookup which didn't consider the trunked ports scenario in L2FW configuration. This issue is fixed in 4.2.2 release by including the trunked ports while doing the FDB lookup.

Q01747321 With certain fctl settings on the dual ports (ports 3-6), some times port status on the backup accelerator are shown as "Link up but Blocking" mode. Although it's not consistent, issue was also seen by constantly changing the various port configurations and alternating between the preferred & backup settings. The issue is caused due to invalid updation of dual port's flag status by MP. This issue is applicable for SFA model 6600 only.
This issue is fixed in 4.2.2 release by correctly updating the flag status with the configured port settings.

Q01778659 When NSF is configured with a static default route and it also learns a duplicate default route via OSPF through an external LSA, the preference will be given the statically configured default route. Now if static default gateway is disabled, the dynamically learned default route is not applied to the kernel routing table and the traffic is affected. This issue is fixed in 4.2.2 by reapplying the OSPF default route in the kernel routing table after the static default route is disabled

Q01820927 In a HA setup with 2 6400 SFA's acting as master & backup, the backup SFA sends outgoing network traffic with multicast source mac. This behavior can become a problem when the accelerators are connected via an intermediate 8600 or core switch may not forward the packets/frames because the source mac is a multicast mac. When the packet is sent out from SFD, the source mac of the packet is mangled with port mask of outgoing interface. When the outgoing interface has port 25 associated with it, then the source mac becomes multicast mac. This issue is fixed in 4.2.2 release by modifying the source mac calculation for outgoing traffic.

Q01832620 Under stress conditions with heavy traffic load, any link flap of the sync port will cause the sync port to go down and it requires a reboot to recover the port. Heavy traffic would mean lot of sync traffic flowing between the cluster members. And when the sync link goes down, all the sync traffic is queued in the transmit buffer of the particular port. The underlying network driver's watchdog continuously monitors the queue and if it detects the queue being full, it'll reset the queue which means reinitializing all the port registers. This issue is fixed in 4.2.2 release by adding a check to see when the transmit buffer is full. If the link status during that time, transmit buffer queue is reset.

Q01809464 While sending responses to SNMP GET requests, the source address of the replies would be chosen based on the outgoing interface IP address even if the request was initially sent to the MIP. To address this issue, a new set of CLI commands /cfg/sys/adm/snmp/adv/getsrcip are added in 4.2.2 release. Users can now configure which IP address to be used while sending the responses to SNMP GET requests

Q01300144 A new BBI enhancement is added in release 4.2.2 to display link speed of SFA ports

8. New Outstanding Issues

Q01863867 Release 4.1.4.1 and above includes new drivers to support RoHS compliant BCM5823 VPN accelerator cards. However, the new drivers are not fully compatible with the older BCM5822 cards. If you're using a version above 4.1.4.0 and has BCM5823 card and are experiencing any traffic interruptions, consider changing the encryption settings to DES-SHA1 or 3DES-MD5.

Q01871072 When a static ARP entry is added and a same ARP entry is also learned dynamically from a duplicate host, the traffic may still be redirected to the learned static ARP there by affecting the traffic. As a work around, make sure no duplicate ARP entries are added via CLI and there're no duplicate hosts in the network

Q01880948 When multiple gateways are configured in NSF, the following issues occur.

1. When more than 1 def gateways are configured, and when metric is changed, duplicate gateways are getting added in the Accelerator's routing table.
2. Default gateways is getting duplicated in the isd routing table
3. Traffic is not getting dropped even after disabling all the def gateways.

Q01877521 If the status of the NAT is changed from no NAT to NAT enabled or from with NAT to without NAT and a CP policy is pushed then no more calls can be made.

Q01877515 Unable to make Voice calls from external network to an internal network. For an attempt to do so the result is a warning stating "Authentication Required"

Q01877961 CLICBD is crashing for fifth and sixth SFDs in a 6 SFD Cluster environment. Following are the two issues that are noticed as a result of the above crash.

1. MIP is not migrating to the fifth and sixth SFDs when we halt the rest of the four SFDs.
2. Hard rebooting the Master SFA will crash the CLI of the fifth and sixth SFDs and the CLI never comes back until we boot delete both of them and rejoin again.

Q01877967 CLI of the SFD will be crashed when we change the type of the SFD from MASTER to BACKUP.

Q01877202 Dynamic routing commands are not working in 5th and 6th SFD connected in a cluster.

Q01881930 Hitless upgrade from 4.1.2.0c_R55 to 4.1.5_R60 failed in a cluster setup with 4 SFD's and 2 SFA's. Part of the cluster has upgraded to 4.1.5 while the other part of the cluster stayed at 4.1.2.

Q01882986 Hitless upgrade from 4.1.2_R60 to 4.2.2_R65 failed in a cluster setup with 2 SFD's and 2 SFA's. Hitless upgrade was successful when upgrading from post 4.1.5 releases to the 4.2.2 release.

Q01879405 While using SIP specific policies with NAT enabled on the Firewall running R65 feature pack, multiple issues were seen. Users can dial from a host in the external network to a host in the internal network, but cannot talk to each other.

In another scenario, after leaving the connection idle for more than 10 minutes, dialing any external phone will fail and a reset of the phone is required.

SIP with no-NAT configuration works fine.

9. New Known Limitations

Hitless upgrade between major Check Point feature packs is not supported. For example, hitless upgrade between R55 and R6x is not supported. Traffic may not be fully synchronized between the cluster members, which can result in some traffic loss. Use the regular upgrade method as described in the upgrade section above while upgrading from R5x to R6x Check Point versions.

For information on previous releases, please refer to the 4.2.1 ReadMe file.

For other known issues, please refer to the product release notes and technical documentation available from the Nortel Technical Support web site at: http://www.nortel.com/support.