



1. Release Summary

Release Date : October 2009
Purpose : Software maintenance release to address customer software issues

2. Important Notes Before Upgrading to This Release

Upgrade to the current release is now limited to support releases from past 18 months. That is, upgrade to 4.2.6 is supported from 4.1.6 or later releases only.

Customers running pre-4.1.6 release and would like to upgrade to 4.2.6 are recommended to do incremental upgrades. In other words, first upgrade to 4.1.6 or later release and then upgrade to the 4.2.6 release.

4.2.6 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”.

Upgrade procedure is the same as mentioned under “[Procedure to upgrade from CLI](#)” section in 4.2.2 Readme section except that the only be upgrade to 4.2.6 from 4.1.6 or above releases only.

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 “/boot/software/hitless/activate” command from 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** is the same as mentioned under “[Procedure for HITLESS upgrade from CLI](#)” chapter in 4.2.2 Readme section.

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 4.2.6 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 4.2.6 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_4.2.6.0_R65.iso (based on R65 + HFA-50)	.iso (contains .img and .pkg files)	486048

File Name	MD5 Checksum
NSF_Director_4.2.6.0_R65.iso	65128009b1714d72b940d52d2313b4fa

5. Version of Previous Release

Software version 4.2.5, Release date - April 2009

6. Compatibility

N/A

7. Changes in This Release

Problems resolved in this release: 4.2.6

Q02015234 In a cluster setup, a passport 8600 is configured as the default gateway to the NSF and is connected via trunk consisting of 4 ports. Whenever one of the trunk ports go down, the gateway status starts flapping which causes the network outage. The gateway status change also causes the configuration to be reapplied which would increase the local_state registry file. And that eventually requires all the services to be restarted on the Firewall which would mean more outage.

The issue is caused due to deletion of ARP entry for the gateway whenever the trunk port goes down. A condition is added in the 4.2.6 to check if the port that's going down is part of a trunk group. If yes, then the trunk status instead of individual port status is used to determine whether to flush the ARP entries or not.

Q02064144-01 When performing SNMP operations (get/getnext/getbulk/walk) on the asfArpTable (OID 1.3.6.1.4.1.1872.2.3.3.1.3.3), socket errors were reported by the CFGD module which causes the entire configuration to be reapplied. This cause the

configuration file to reach the maximum file size limit after which the services have to be restarted manually that can cause traffic impact.

The issue is caused by failure to close the socket used to get the SFA MAC entries to fetch the Arp table dump. Code has been modified in 4.2.6 to use the local registry file for the SFA Mac address information instead of using sockets.

Q02009236

/maint/schedule command which is used to schedule automatic uploading of exdumps used for debug/monitoring purposes does not clear the files after pushing the logs to the remote location. This causes the /var/ partition to reach the maximum limit after which cause the cpu to spike.

Code has been corrected in 4.2.6 to delete the backup files once the file transfer is complete.

Q01874246-02

In a cluster with more than 3 or more Directors and sync enabled, each Director crashes randomly with the panic signature pointing to “acct_info_request_cb” function.

The issue has been addressed in Check Point’s R65 HFA50 which is integrated into the 4.2.6 release.



1. Release Summary

Release Date : April 2009
Purpose : Software maintenance release to address customer software issues

2. Important Notes Before Upgrading to This Release

Upgrade to the current release is now limited to support releases from past 18 months. That is, upgrade to 4.2.5 is supported from 4.1.4 or later releases only.

Customers running pre-4.1.4 release and would like to upgrade to 4.2.5 are recommended to do incremental upgrades. In other words, first upgrade to 4.1.4 or later release and then upgrade to the 4.2.5 release.

4.2.5 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”.

Upgrade procedure is the same as mentioned under “[Procedure to upgrade from CLI](#)” section in 4.2.2 Readme section except that the only be upgrade to 4.2.5 from 4.1.4 or above releases only.

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 “/boot/software/hitless/activate” command from 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** is the same as mentioned under “[Procedure for HITLESS upgrade from CLI](#)” chapter in 4.2.2 Readme section.

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 4.2.5 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 4.2.5 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_4.2.5.0_R65.iso (based on R65 + HFA-30)	.iso (contains .img and .pkg files)	464576

File Name	MD5 Checksum
NSF_Director_4.2.5.0_R65.iso	a5090ef50baa7153f80ffe1ca983e273

5. Version of Previous Release

Software version 4.2.4, Release date - October 2008

6. Compatibility

N/A

7. Changes in This Release

Problems resolved in this release: 4.2.5

- Q01945468 Hitless upgrade fails when upgraded from prior to 4.1.4 builds, if a port is configured with preferred fiber, backup copper, flow control none and auto-negotiation is disabled options on the copper interface.
This is due to new validation added in 4.1.4 when autoneg is disabled. This issue is fixed in 4.2.5 release.
- Q01231858 Sometimes attempting to browsing websites is failing. Analysis showed that the issue due to incorrect assumptions while implementing TCP sequence verification (TSV) feature.
TSV feature is corrected as per Check Point's SecureXL design and the issue is fixed in 4.2.5 release.
- Q01961000 Policy push reduces free space in the local state registry and if it filled up to 90% this is causing all the services to restart.
Code has been modified such that an alarm/trap will be raised once the free space

in local_state reduces to 20% of its size. The alarm message would also give the user a work around to limit the impact due to restart of services by manually restarting the services.

This enhancement is added in 4.2.5 release.

Q01896498 Unnecessary Packet drops are seen for Generic FTP traffic, “fw ctl kdbg drop” shows many packet drops as reason "INVALID ACK".

TSV feature is corrected as per Check Point’s SecureXL design and the issue is fixed in 4.2.5 release.

Q01933786 Irrelevant messages are thrown while creating Backup from CLI. This is due to the 'tar' command throwing warning messages for non existent files.

This issue is fixed in 4.2.5 release by modifying the code not to throw any irrelevant messages.

8. New Outstanding issues

N/A

9. New Known Limitations

N/A

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>.

Copyright © 2009 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globemark, and Alteon are trademarks of Nortel Networks Limited.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at <http://www.nortel.com/support>.



NORTEL SWITCHED FIREWALL 6000 Series Software Release 4.2.4

1. Release Summary

Release Date : October 2008

Purpose : Software maintenance release to address customer software issues

2. Important Notes Before Upgrading to This Release

Upgrade to the current release is now limited to support releases from past 18 months. That is, upgrade to 4.2.4 is supported from 4.1.4 or later releases only.

Customers running pre-4.1.4 release and would like to upgrade to 4.2.4 are recommended to do incremental upgrades. In other words, first upgrade to 4.1.4 or later release and then upgrade to the 4.2.4 release.

4.2.4 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”.

Upgrade procedure is the same as mentioned under “[Procedure to upgrade from CLI](#)” section in 4.2.2 Readme section except that the only be upgrade to 4.2.4 from 4.1.4 or above releases only.

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 “/boot/software/hitless/activate” command from 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** is the same as mentioned under “[Procedure for HITLESS upgrade from CLI](#)” chapter in 4.2.2 Readme section.

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 4.2.4 supports the following *Check Point applications:

- Firewall-1®
- VPN-1®
- SmartDefense™
- NAT
- Authentication
- Content Security
- Policy Server
- Management Tools
- SmartView Monitor™

1

Configure the following management tools outside the NSF 4.2.4 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_4.2.4.0_R60.iso (based on R60 + HFA-06)	.iso (contains .img and .pkg files)	355663872
NSF_Director_4.2.4.0_R65.iso (based on R65 + HFA-02)	.iso (contains .img and .pkg files)	486965248

File Name	MD5 Checksum
NSF_Director_4.2.4.0_R60.iso	b5dbca2f7456f9096588e540368236f5
NSF_Director_4.2.4.0_R65.iso	2512c87bc3c5fd0ba71076f954f645da

5. Version of Previous Release

Software version 4.2.3, Release date - August 2008

6. Compatibility

N/A

7. Changes in This Release

Problems resolved in this release: 4.2.4

- Q01918557 In a dynamic routing environment, when 2 routes (one learned via OSPF and the other a static route) exist to the same network, dynamic route should be deleted and the static route should be added.
Due to the invalid handling of these routes, the order in which these 2 routes are added determines which type of route would get preference.
The delete route function uses the network address & mask to delete the dynamic route. Since the static route also has the same network address & mask, its also getting deleted.
- This issue is resolved in the release 4.2.4 by properly checking for the correct route to delete.
- Q01880948 Duplicate entries of static default gateway are seen in accelerator when the default gateway metric is changed from round-robin to strict.
The default gateway handling routines didn't consider multiple gateways scenario while adding/deleting the default gateway configuration.
- This issue is now fixed in 4.2.4.
- Q01912932 When a port goes down, the corresponding interfaces and default gateways bound to that port should go down and when the port is brought up they should come up. But the default gateways are not coming up even when the port is brought up.
- This issue is solved with fix for CR # Q01880948 available in 4.2.4.
- Q01912913 After enabling & disabling the multiple default gateways, SFA routing table should display only the enabled default gateways. Due to invalid handling of the delete route function as specified in CR # Q01880948, entries of disabled default gateways are also seen in the accelerator routing table.
- This issue is resolved in 4.2.4 by properly handling all the configured gateways.
- Q01892695 All the existing connections were lost when any new configuration is applied in NSF when only sync2 is enabled. This issue doesn't occur when only sync1 or both sync1 and sync2 are enabled.
Whenever any new configuration is applied, there's a check to see if any sync configuration is changed. If yes, all the CP services need to be restarted. However, this condition was always returning true when only sync2 is enabled thus restarting CP services every time a new config is applied thereby terminating all the existing connections.
- This issue is resolved in 4.2.4 release by properly introducing a check to restart CP services only when any sync configuration is modified.
- Q01882986 Hitless upgrade fails when upgrading from pre-4.1.5 version to any post-4.1.5 releases. The design of hitless upgrade is to upgrade one of the SFDs first while the other handles all the sessions. After upgrade, the same SFD will try to upgrade the accelerator connected to it and push the new configuration to accelerators. The

upgraded SFD will yank all the sessions to it and then starts upgrading the remaining SFDs.

When the upgraded SFD pushes the config to both the accelerators, it doesn't check whether the accelerator is upgraded or not.

It may be possible to introduce new configuration support in the upgraded release. And the SFD will try to push the new configuration to the accelerator which is not upgraded yet. This results in the hitless upgrade to fail.

This issue is resolved in 4.2.4 by adding a check to see if the accelerator is upgraded or not, only then push the new configuration.

Q01896432 Hitless upgrade fails when upgraded from 4.1.5 to 4.2.3.
This issue is also similar to CR # Q01882986 in the sense that a new CLI configuration option “/cfg/vrrp/fastfail” which was added in 4.2.3 release causes the hitless upgrade to fail.

This issue is resolved in 4.2.4 with the fix given for CR # Q01882986.

Q01921779 No OSPF HELLO packets are coming out of port 25 in NSF setup with 6400 accelerators.
While sending the packets out from the SFD, the source MAC is calculated to include the outgoing port & VLAN number for SFA to use while sending the packet of the correct port.
This behavior was slightly changed in the previous release as part of CR # Q01820927 that caused the current issue.

The issue is resolved in the release 4.2.4 by modifying the source MAC calculation to handle port 25 properly.

8. New Outstanding Issues

Q01938302 When /var drive is filled up completely with logs and hard disk usage shows 100%, cfgd and tngsys logs stop normal logging as expected. The logging doesn't recover even after manually freeing the disk space of /var.

Work around is to restart the cfgd and syslog-ng services for the logging to recover.

Q01892073 A new vulnerability US-CERT VU#878044 is identified that allows attackers to read and modify any SNMP object that can be accessed by the impersonated user. Attackers exploiting this vulnerability can view and modify the configuration of these devices.
NSF products are currently vulnerable to this attack and it's recommended to turn off SNMPv3 as a short term work around.

9. New Known Limitations

- Q01930476 When port mirroring is enabled on any of the ports of the master accelerator then VRRP fail over is happening. This issue is due to the design of VRRP in which all the enabled ports (irrespective of whether there are mirrored ports or not) are used for calculating priority and hence any change in the link status on any of the enabled ports on master accelerator invokes VRRP failover process.
- Q01945468 Hitless upgrade could fail when certain configuration settings are not supported on the upgrade version. For example, a new validation was added in 4.1.5 release which doesn't allow configuring a port with autoneg OFF and backup as NONE. When the user tries to upgrade with the above setting to 4.1.5 or above release, the hitless upgrade process may fail.
This behavior is not restricted to the above validation alone and can occur due to any new validation added in the particular release.

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>.

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globemark, and Alteon are trademarks of Nortel Networks Limited.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at <http://www.nortel.com/support>.



NORTEL SWITCHED FIREWALL 6000 Series

Software Release 4.2.3

1. Release Summary

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

2. Important Notes Before Upgrading to This Release

Upgrade to 4.2.3 is supported from 4.1.1 or later versions. 4.2.3 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 4.2.3, please keep the following things in mind. 4.2.3 is a combined L2/L3 firewall. If you have multiple ports in the same VLAN, the default behavior of 4.2.3 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.

For information on CRs # Q01158579 and Q01157140 please refer to 4.1.x Read Me.

Upgrade procedure is the same as mentioned under “[Procedure to upgrade from CLI](#)” section in 4.2.2 Readme section.

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 “/boot/software/hitless/activate” command from 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** is the same as mentioned under “[Procedure for HITLESS upgrade from CLI](#)” section in 4.2.2 Read Me.

Note: *Hitless upgrade will fail from 4.2.1 or previous releases to 4.2.3 release.*

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 4.2.3 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 4.2.3 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_4.2.3.0_R60.iso (based on R60 + HFA-06)	.iso (contains .img and .pkg files)	355663872
NSF_Director_4.2.3.0_R65.iso (based on R65 + HFA_02)	.iso (contains .img and .pkg files)	486965248

File Name	MD5 Checksum
NSF_Director_4.2.3.0_R60.iso	addeb8c04f4a4fab92d5796720a19b19
NSF_Director_4.2.3.0_R65.iso	4454600114eb0485b1135c0c47c2c3a8

5. Version of Previous Release

Software version 4.2.2, Release date - June 2008

6. Compatibility

N/A

7. Changes in This Release

Problems resolved in this release: 4.2.3

- Q01865541 Link goes down when the speed of the dual ports is set to 1000Mbps with the following configuration – “Autoneg” is ON, “preferred” set to copper and “backup” set as fiber. The same issue does not occur when the preferred is set to fiber and the backup to copper. Also the issue is not seen with dedicated fiber ports. The issue is resolved in the release 4.2.3 by not using the manual port settings when autoneg is set to ON.
- Q01866402 NSF allows a maximum of 6 SFDs in a cluster. As per the design, first four SFDs would be treated as ‘Master’ while the remaining two would be treated as “Slaves”.

Slave SFDs cannot become MIP in case the MIP owner is down/not reachable. They always wait for a MIP to start the services. This is done so as to limit the MIP election process in case of a failure.

All the cluster members talk over a dedicated port which is taken outside the purview of Check Point rule base. SSI bypass feature takes care of handling the communication between the cluster members.

However, during the MIP fail-over process, the slave SFDs send broadcast packets to enquire about another MIP owner in the network. And the SSI bypass feature was not handling the broadcast traffic. So this traffic is handled by Check Point and if there's no explicit rule to allow it, the packets would be dropped.

As the CP module is dropping these broadcast packets, the slave SFDs are unable to reach MIP owner, which results in its failure to join the cluster.

The issue is solved in the release 4.2.3 by letting these broadcast packets handled properly without the need for any explicit Check Point rule.

Q01870563 Accelerator goes to ACCEL-OFF state when more than 128 static arp entries are added. Accelerator can only support up to 128 static arp entries, but there was no validation on the SFD not to allow more than 128 entries. When the faulty configuration is pushed to the accelerator, it results in accel off state. This issue is resolved in the release 4.2.3 by adding a validation on the maximum number of added static arps.

Q01871072 Consider a network with two hosts (host1 and host2) with same IP address. When traffic flows from either of these hosts across the NSF, accelerator stores an ARP entry with host1 Ip address and host2 mac address. If a user now adds a static ARP entry for the same host1 Ip address but with host2's MAC address, configured static MAC address should get preference while forwarding the traffic. But traffic on host1 continues and host2 cannot be reached. This is caused due to an error in the static arp handling. The issue is fixed in the release 4.2.3 by properly updating the SP arp cache with the added static arp.

Q01871983 The /var/tmp/sensors file does not rotate after reaching the max file limit. The issue is resolved in the release 4.2.3.

Q01872414 Upgrade from BBI to 4.2.2_R65 from any lower version fails since the newly added Check Point's HFA-02 for R65 has increased the package size to more than 200MB. The maximum upload file size from BBI is set to 160MB, which is far less than the actual size of the 4.2.2 package. This issue is resolved in the release 4.2.3 by setting the maximum allowable size of the package to 300MB from previous 160 MB.

Q01872471 The existing FTP connections are getting terminated when any FIN packet with

invalid sequence number intruded into the firewall. This issue is caused because of a bug in the code, which was corrected as part of an enhancement to the TCP sequence Verification issue and it is resolved in 4.2.3 release.

Q01847032 Port fctl settings are set properly when autoneg is ON. When autoneg is ON, manual port settings should not be considered while applying the configuration on the accelerator ports. But there was no check to prevent the manual port settings being applied on the SFA with autoneg ON. Due to this, the fctl settings are not properly set as expected.

The issue is fixed by adding a check in SFD for the status of autoneg before updating the port configurations to the accelerator. When autoneg is ON, no port configurations will be updated to the accelerator and will be auto negotiated.

Q01877961 In a cluster of 6 SFDs connected to 2 SFAs, first 4 SFDs would be automatically configured as 'Master' and the remaining two would be configured as 'Slave'. Only Master SFDs can take MIP ownership in case of a failure. This is designed so as to limit the MIP election procedure.

Rebooting the master SFA will result in the Slave SFDs losing the connectivity with the MIP. All the CLI operations on the Slave would fail due to this. But even after the connectivity is restored, the Slave SFDs cannot recover.

The reason for this issue is similar to CR # Q01866402 and fixed in 4.2.3.

Q01831671 Route learnt via type-3 LSA is not added to route table if the same route learnt via type-5 LSA is available. After initially adding a type-3 route, any similar type-5 is not added to route table. OSPF then sends a message to delete this type-5 route without specifying the gateway. Since type3 route is also same as the type-5 route, this route it gets deleted.

This issue is fixed in 4.2.3 by correcting the logic of deleting the type-3 and type-5 route entries.

Q01863867 Traffic is affected when BCM5822 with CPacc4 module and BCM5823 with CPacc3 module are installed. This issue is resolved in the release 4.2.3 by modifying the install script, which will load the appropriate CPacc3 module for BCM5822 and CPacc4 for BCM5823 VPN accelerator cards.

The fix for this CR has also fixed the following CRs:

- Q01870058 - In secure client mode of VPN setup, the remote secure client PC is able to connect to its gateway but the same is not able to get any desktop security rules and not able to send the traffic.
- Q01870063- In secure remote mode of VPN setup, the secure remote PC is able to connect to its gateway but the same is not able to handle any traffic.

Q01855791 L2FW over trunked ports doesn't work as desired and regular traffic is affected. This issue is caused due to incorrect assumptions in the FDB lookup. FDB lookup returns trunk ID as the egress port, which is treated as a port number for the L2FW processing.

This issue is resolved by checking the return type from the FDB lookup. If it returns a trunk Id, a hash is calculated on all the trunk ports upon which the egress port is selected.

- Q01850786 From 4.1.6 onwards, users can configure redundant sync interface. However, the default 2nd sync interface and the management interface were configured to use the same physical device. Due to this, management interface configuration is not getting applied if 2nd sync is not configured. Also the management interface settings are over written with the 2nd sync settings.
The issue is resolved in the release 4.2.3 by correcting the default devices for MGMT device and Sync2.
- Q01856573 During a policy push, acceleration would be turned OFF and turned ON again. This is required to apply any changes in the rule base to the existing & new traffic. As per the original design, during accel OFF, the particular SFD would be marked for deletion which would mean that it cannot handle the traffic.
Hence when the SFA receives data that needs to be forwarded to a particular SFD that's marked for deletion, the session would be re-bound to the next available SFD.
This behavior works well when there is a synchronization of that particular session among the SFDs. But synchronization is not done for all the services, for eg. It is recommended by NORTEL to turn OFF sync for services like http. In this case the behavior stated would cause the termination of the sessions, which is not an expected behavior.
The issue is resolved in 4.2.3 by modifying the SFA behavior during policy push. SFA would now forward the traffic to the same SFD which was handling the traffic before "ACCEL OFF" so that the traffic won't be dropped due to non-synchronization.
- Q01794609 The System LED description as given in the Hardware Installation Manual is incorrect.
The system status LEDs indicates the operational status of four fans, chassis, CPU temperature, ambient temperature and the voltages (+5V and +12V).
The different glow states of LEDs are as follows:
- If the system is reset, the LED doesn't glow
 - If system detects any problem with any of the CPU temperature, fan speed or system voltages, the LED glows amber
 - If the system is working, the LED glows in solid green
 - If the system halts, LED flashes

8. New Outstanding Issues

- Q01612783 TCP connections with TCP window-scaling option enabled, stall intermittently when the session is started. The problem is found to be with CP firewall dropping the initial packets due to TCP sequence error. These dropped packets when retransmitted by server continue the service there after. The same problem doesn't occur.
- Q01896432 Hitless upgrade fails from the version 4.1.5_R65 to 4.2.3_R65. Normal upgrade

works fine in this case. So users upgrading from 4.1.5_R65 to 4.2.3_R65 are requested to use the normal upgrade instead of Hitless Upgrade. This issue is planned to fix by next release.

9. New Known Limitations

Hitless upgrade will fail from 4.2.1 or previous releases to 4.2.3 release. So users are requested to use normal upgrades for this particular case.

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>.

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globemark, and Alteon are trademarks of Nortel Networks Limited.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical Support on the web at <http://www.nortel.com/support>.



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 4.0.1 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/cw” 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*” command. 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 “/boot/software/hitless/activate” command from 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.

9. New Known Limitations

N/A

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>.

Copyright © 2007 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, Globemark, and Alteon are trademarks of Nortel Networks Limited.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel. To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Technical

Support on the web at <http://www.nortel.com/support>