Ethernet Routing Switch 5510/5520/5530

Software Release 5.1.2

1. Release Summary

Release Date: 30-September-2008
Purpose: Software patch release to address customer found software issues.

2. Important Notes Before Upgrading to This Release

For customers upgrading from older software versions, a series of upgrades are required to prevent configuration corruption under certain circumstances. This upgrade path includes the following releases: 4.0, 4.1, 4.2 and 5.0.

3. Platforms Supported

Ethernet Routing Switch 5510/5520/5530

4. Notes for Upgrade

Please see "System Configuration Guide for Nortel Ethernet Routing Switch 55xx Series, Software Release 5.1" (Part No. 217468-B, available at http://www.nortel.com/support). In the Product Finder, select Routers and Routing Switches- Ethernet Routing Switch 5510, 5520, or 5530-24TFD, followed by Documentation) for details on how to upgrade your Ethernet Routing Switch.

File Names for This Release

File Name	Module or File Type	File Size (bytes)
5530_500004_diag.bin	Diagnostic image	812,036
5530_512034.img	Agent code image	6,000,520
5530_512035s.img	Agent code image (SSH)	6,244,724

5. Version of Previous Release

Software Version 5.1.1

6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.9.

7. Changes in This Release

7.1. New Features in This Release

7.1.1. Dynamic VLAN assignment from RADIUS server for EAP and non-EAP authentication device

In this release both EAP and non-EAP clients may be authenticated and assigned to a VLAN, when in MHMA mode.
In other modes, VLAN assignment is ignored.
In order to utilize this feature, the user must enable the radius-mac-authentication globally and at the interface (port) level.
Once enabled, EAP and non-EAP MACs will be assigned to configured VLANs.

CLI commands

New CLI commands are provided for Radius assigned VLANs for EAP and Non-EAP MACs:

```
EAPol multihost use-radius-assigned-VLAN
no EAPol multihost use-radius-assigned-VLAN
default EAPol multihost use-radius-assigned-VLAN
EAPol multihost non-EAP-use-radius-assigned-VLAN
no EAPol multihost non-EAP-use-radius-assigned-VLAN
default EAPol multihost non-EAP-use-radius-assigned-VLAN
```

7.1.2. Radius Assigned VLAN Update for 802.1x - Use most recent Radius VLAN Enhancement

Existing Functionality
If use-radius-assigned-vlan option is enabled, the first valid radius-assigned-vlan (by EAP or Non-EAP authentication) on that port will be honored. Subsequent radius-vlan assignments will be ignored, for any user on that port. Note: If EAP VLAN is assigned after the Non-EAP VLAN, then the EAP VLAN takes precedence over the non-EAP with the port being moved from the Non-EAP VLAN to the EAP VLAN.

New Functionality (as implemented in 5.1.2 release)
The new functionality introduced with release 5.1.2 is to honor the last received radius-vlan assignments on a port. The last radius-assigned VLAN (either EAP or Non-EAP) will determine the VLAN membership and PVID replacing any previous radius-assigned VLAN values for that port.

Functional examples:
1. NEAP (Non-EAP) device authenticates on port X, the VLAN membership and PVID will be changed according to radius-vlan assignment if any. For example, let's assume that it will be added to VLAN 50.
2. If a PC authenticates on the same port X, that port will be removed from VLAN 50 and moved into the new radius-vlan assigned with the new PVID (if a VLAN is assigned from the radius server). The VLAN and the PVID will be changed every time a new valid radius-vlan assignment is processed on the port as a result of a new authentication.

Other functional example:
1. Multiple EAP and NEAP (non-EAP) clients authenticate on a port.
2. The EAP clients perform re-authentication; the non-EAP clients age out and are re-authenticated; the last VLAN assigned setting for either EAP or NEAP clients will always be applied (meaning there is a potential for the VLAN to swap based on re-authentication).
CLI, SNMP and ACG interfaces are supported (no WebUI is available for this function)

CLI commands

```
eap multihost use-most-recent-radius-vlan enable

no eap multihost use-most-recent-radius-vlan enable
```

ACG support will be represented by the presence of either of the following lines in the EAP configuration section of the ASCII:

```
eap multihost use-most-recent-radius-vlan enable
```

or

```
no eap multihost use-most-recent-radius-vlan enable
```

SNMP Support

Global:

```
bseeMultiHostUseMostRecentRadiusAssignedVlan OBJECT-TYPE
        SYNTAX TruthValue
        MAX-ACCESS read-write
        STATUS current
        DESCRIPTION "Controls whether to use most recent RADIUS-assigned VLAN."
        DEFVAL { false }
        ::= { bseeObjects 21 }
```

Per Interface:

```
bseePortConfigMultiHostUseMostRecentRadiusAssignedVlan OBJECT-TYPE
        SYNTAX TruthValue
        MAX-ACCESS read-write
        STATUS current
        DESCRIPTION "Controls whether to use most recent RADIUS-assigned VLAN."
        DEFVAL { false }
        ::= { bseePortConfigEntry 18 }
```

7.1.3. CPU utilization and Memory Utilization

This feature provides the CPU utilization and memory utilization of all the units in the stack. The CLI command "show cpu-utilization" and a MIB s5ChasUtilTable in the S5-CHASSIS-MIB, are added for this feature.

The "**show cpu-utilization**" command displays the CPU utilization of individual units in a stack as shown below:

```
5520-48T-PWR(config)#show cpu-utilization
----------------------------------------------------------------
      CPU Utilization
----------------------------------------------------------------
Unit/ Last 10 Sec, 1 Min, 10 Min, 60 Min, 24 Hrs, System Boot-Up
----------------------------------------------------------------
1          25%      25%   24%     NA       NA      26%
2          24%      24%   24%     NA       NA      25%
```

The CPU utilization displays as a percentage of the CPU load on the specific unit. The values show the CPU utilization for the last 10 seconds, 1, 10 and 60 minutes, 24 hours and since system startup.
Normal values for the CPU utilization are less than 30% CPU load. In a stack configuration this value is around 30-35% CPU load depending on stack size. When the switch is loaded for a period of time with some kind of processing (e.g. show running-config command) the CPU utilization values will increase.

This feature also provides CLI command, “**show memory-utilization**” to display the percentage of dynamic memory (memory left for applications after the switch initialized) that is currently available on each unit in stack and a minimum dynamic memory available mark since unit boot. The output of the memory-utilization command is given below:

```
5520-48T-PWR(config)#show memory-utilization
--------------------------
      Memory Utilization
--------------------------
Unit/ Available Low Mark
--------------------------
1     99%         99%
2     99%         99%
```

After switch/stack startup the memory utilization should be around 99% available. When the switch performs some operations after startup such as image download, the available memory value will decrease. After the operations have completed, the memory will increase to its former value. The low mark will save the least (lowest number) available memory since startup. In normal switch/stack operation there should always be available memory for applications to use.

<u>NOTE</u>: The CPU utilization may vary due to different network configurations. It is important for the customers to note their normal CPU utilization when the network is stable. A high CPU utilization for a prolonged period may be due to problems in the network.

7.1.4. Multicast/Broadcast Storm Control

The purpose of this feature is to provide Broadcast and/or Multicast storm protection. The Multicast/Broadcast storm control has been implemented as part of rate-limiting. The Rate Limiting implementation on earlier release supports rate limiting based on percentage. This release also supports rate limiting based on packets-per-second settings.
Two modes for rate limiting are now available:

percent mode: the user can set the forwarding rate for the specified packet types (broadcast, multicast or both) not to exceed a specified percent (1-10%) of the available bandwidth on a port. The current implementation does not change anything regarding the percent mode.

pps mode: the user can set the forwarding rate for specified packet types (broadcast, multicast or both) not to exceed a specified number of packets per second.

For example, is a user had set the rate limit for broadcast traffic to 10%, using the percent mode, and the bandwidth on that port is 100Mbps, all broadcast traffic that exceeds 10Mbps will be discarded. When using pps mode, if the user had set the limit to 1000 pkts/sec, any additional broadcast packets that exceed the threshold will be discarded.
This feature controls the incoming broadcast/multicast traffic using three parameters: mode, type of packets and threshold.
Mode can be:

Percent: Default

pps

Type can be:

Multicast

Broadcast

Both

Threshold:

0-10: for Percent mode

0-262143: for pps mode

NOTE: The Broadcast type is used to limit *broadcast* and *unknown unicast* traffic.
If pps mode is used, a pkts/sec value can be set directly by the user. The rate value set will go to the rate limit registers and the specified type of traffic will be limited.

Maximum pkts/sec value. The maximum rate the user can set is 262143 pkts/sec.
Speed change events: The pkts/sec rate set will remain unchanged.

Upgrade/downgrade
When upgrading/downgrading to a version that does not support the rate limit with pps mode, ports having rate limit settings in pps mode will take the default values (packet type Both, percent None).

CLI Commands:
The rate limit can be set and displayed for both modes (percent and pps) from CLI.

1. The new CLI command for setting the rate limit is:
rate-limit {both|multicast|broadcast} {percent <0-10> | pps <0-262143>| <0-10>}

For backwards compatibility, the old command that sets the rate-limiting based on a percentage was retained, and a mode option was added to it. If the user uses this command with no mode specified (percent/pps), the rate-limiting will be set based on percentage.

2. The rate limit settings will be displayed according to the mode used, through the following command:
show rate-limit {port <LINE>| <cr>}

Example:

```
5520-48T-PWR(config)# interface FastEthernet 1
5520-48T-PWR(config-if)#rate-limit broadcast percent 2
5520-48T-PWR(config-if)#interface FastEthernet 2
5520-48T-PWR(config-if)#rate-limit multicast pps 12000
5520-48T-PWR(config-if)#show rate-limit port 1-2
Port Packet Type   Limit        Last 5 Minutes Last Hour Last 24 Hours
---- ----------- -------------- -------------- --------- -------------
 1    Broadcast   2%                0.0%        0.0%        0.0%
 2    Multicast   12000pkts/sec     0.0%        0.0%        0.0%
```

3. To disable rate limiting settings:
no rate-limit

4. The default rate limit settings are mode *percent*, type *both*, percent value *None*:
default rate-limit

5. Setting rate limit threshold 0:

In both modes (percent and pps) the rate limit threshold can take the value 0, effectively disabling rate limiting on the specified port:
rate-limit {both|multicast|broadcast} 0
rate-limit (both|multicast|broadcast) pps 0
Please note, that this is not the equivalent of *default rate-limit*. If entering one of the two commands, there will be no filtering on the specified interfaces, but the entered mode (percent/pps) and the type (broadcast/multicast/both) will be set.

NOTE: On 10GIG ports, the minimum threshold that can be set in pps mode is 1000 pkts/sec. If the user tries to set a lower value, an error message will be displayed: "The threshold value has to be at least 1000 on 10GIG ports". The exception to this rule is value 0, which can be set on 10GIG ports also.

Console Interface

The same menu as previous releases is used: Switch Configuration Menu -> Rate Limiting Configuration.
percent mode. The exact configuration and display characteristics as in previous releases are available, backward compatibility being preserved.
pps mode. The Limit column can take a new value, "PPS", which is only used for displaying purposes. For ports with rate limit settings in pps mode, the CI will show the type set in the Packet Type column, and "PPS" indicator in the Limit column.
Having the example given in the CLI section, this is how the CI display will look like:

```
                           Rate Limiting Configuration
Port Packet Type   Limit   Last 5 Minutes Last Hour Last 24 Hours
---- ------------- ------- -------------- --------- -------------
1    [ Broadcast ] [ 2% ]   0.0%           0.0%       0.0%
2    [ Multicast ] [ PPS ]  0.0%           0.0%       0.0%
```

Port 2 has rate limit set in pps mode, type multicast packets, value 12000pkts/sec (indicated with PPS).
Setting the rate limit in PPS mode is only available from CLI. The PPS option is skipped when navigating through the CI Limit column.

WEB/JDM

Web/JDM have only the percent mode available. For percent mode, the exact configuration and display, as in previous releases, is used. For ports having rate limit settings in pps mode, the type set (Broadcast, Multicast, Both) will be displayed and percent "None".
NOTE: If the user sets the rate limit for type "Both", using pps mode, the Web/JDM interfaces will display type "Both" and percent "None", these values being the same with the default settings. For security, if using the pps mode, CLI and CI should be used for displaying the rate limit settings.

SNMP

The PPS enhancement has no support for SNMP. The percent mode keeps the old SNMP configuration available, with no changes.

Limitation on 10GIG ports

Because of a hardware limitation, the rate limit PPS mode for 10 GIG ports has a restriction for setting the pkts/sec threshold. The minimum rate that can be set on these ports is 1000 pkts/sec.
The rate limiting in pps mode on these ports works in approximate 1041 ranges.
Below are some examples of the threshold the user can set and the actual results of rate limiting:
1000 - 1041 pkts/sec ----> the real rate limit threshold will be around 1041 pkts/sec.
1042 - 2083 pkts/sec ----> the real rate limit threshold will be around 2083 pkts/sec.
2084 - 3125 pkts/sec ----> the real rate limit threshold will be around 3125 pkts/sec.
3126 - 4167 pkts/sec ----> the real rate limit threshold will be around 4167 pkts/sec.
and, so on.

7.1.5. Automatic Unit Replacement manual disable

The AUR feature currently provides the following functionality:

a. Saving Configuration

The function will automatically save the configuration of all non-base units to the base unit. When the configuration of a non-base unit is updated, the configuration of this unit is **automatically** sent (synced) to and saved in the base unit.

b. Restoring Configuration

This function will **automatically** restore a saved configuration to a new non-base unit (different MAC) in the stack. This restoring function can be enabled/disabled via CLI command. The default is enabled.

Enhancement:

In this release, the following enhancement is provided:

User is able to enable/disable the configuration-saving function via CLI commands. The enable/disable state of this function is retained across a reset. The default mode is enabled.

User is able to manually restore an associated configuration (same Unit Number) to a non-base unit (regardless of MAC).

User is able to manually save a configuration of a non base unit to the base unit regardless the state of the AUR feature.

CLI commands:

1. Enable and Disable of Configuration-saving

stack auto-unit-replacement config save enable

This CLI command will re-enable the automatic saving of the configuration of the non-base units.

stack auto-unit-replacement config save disable

This CLI command will stop saving the configuration of the non-base units to the base unit. The last saved configuration of non-base units in the base unit will be retained.

NOTE: **The enable/disable state of this function is saved across a reset**.

2. Manual Restore

stack auto-unit-replacement config restore unit <1-8>

This command will restore an associated configuration (same Unit Number) to a non-base unit n (regardless of MAC).

3. Manual Save

stack auto-unit-replacement config save unit <1-8>

This command will save the configuration of a non base unit (1-8) to the base unit regardless of the state of the "config-save".

NOTE: This command is required to be entered from the base unit console port or via telnet.

4. CLI Command Usage

These CLI commands have to run from the CLI Privileged Mode, not from the configuration Mode.
For example,

```
-> enable /* to get in the CLI privileged Mode */
->## stack auto-unit-replacement config restore unit 2
```

Enhancement for "show stack auto-unit-replacement" command

The command now will provide the following information.

show stack auto-unit-replacement

```
Auto Unit Replacement Auto-Restore: Enabled Auto Unit Replacement Auto-Save: Disabled
Unit # Last Configuration-Save Time-Stamp Ready For Replacement
1 3 days 10:23:02 Yes
2 0 days 00:01:40 No
3 3 days 10:12:33 Yes
6 3 days 10:12:34 No
8 3 days 10:12:35 Yes
```

AUR Behavior:

Auto Unit Replacement Auto-Restore

Enable: During a unit replacement, the configuration will be automatically restored to the new unit.
Disable: During a unit replacement, the configuration will **not** be restored automatically.

Auto Unit Replacement Auto-Save

Enable: The current configuration of a non base unit will be automatically saved to the base unit.
Disable: The current configuration of a non base unit will **not** be automatically saved to the base unit.

Last Configuration-Save Time-Stamp

This system-up time of the non base unit is recorded when the non base unit sends its configuration to the base unit.

Ready for Replacement

Yes: The current configuration of the non base unit has been saved to the base unit. This unit is currently ready for replacement.
No: The current configuration of the non base unit is not saved to the base unit. The latest changes of the configuration of the non base unit will be lost if the unit is replaced with a new unit.

7.1.6. VLAN transition on MAC authentication

VLAN transition on MAC authentication is a feature that requires SNASv2.0 or higher.
With this feature, upon MAC-authentication, the SNAS can request a filter-only change or a VLAN and filter change. Note that VLAN changes would affect all MACs on the port, since we only support port-based VLANS.
No new user configuration is needed at the switch for this feature.

7.1.7. Fail open mode for NSNA ports

The FailOpen feature is used when connection to NSNAS is lost or there is no NSNAS present.
A FailOpen VLAN and filter can be configured at the switch.
They have to be valid NSNA VLAN-IDs in red, yellow or green.
A red fail-open VLAN can be paired with a red, yellow or green filter.
A yellow fail-open VLAN can only be paired with a yellow filter.
A green fail-open VLAN can only be paired with a green filter.

When FailOpen is enabled and NSNAS connection is lost, existing clients will not change VLANs, and network access will not be interrupted for them.
New clients will be moved to the FailOpen VLAN and filter. If the NSNAS connection comes back up, all failOpened ports will be moved back to red, and authenticated by the NSNAS.

7.1.8. TFTP Filename length increase to 128 characters

This enhancement (*Q01654927)* extends the maximum TFTP configuration filename length from 30 to 128 characters. Changes affect all user interfaces: CLI, Console, Web, SNMP and JDM. For now, longer filenames will only be supported for configuration files (both binary and ASCII).

CLI

Filenames of up to 128 characters can now be given as parameters to the "*copy config tftp*", "*copy tftp config*", "*copy running-config*" and "*configure network load-on-boot*" commands. Previously, the filename parameter could have a maximum of 30 characters.

Web and SNMP/JDM

Sets for the following SNMP objects will now be successful with filenames up to 128 characters long: s5AgSysAsciiConfigFilename, s5AgSysBinaryConfigFilename.

Console Interface

Filenames up to 128 characters long can now be entered in the following forms under the "Configuration File" menu:
- "Configuration File Download/Upload" form - "Configuration Image Filename" field
- "ASCII Configuration File Download" form - "ASCII Configuration Filename" field.

Since the Console has some special display limitations, only allowing 80 characters per line, the behavior for these forms will be as follows:
- When entering the filename, if the length of the entered string exceeds 30 characters, scrolling will occur. Thus, the string displayed after the “Enter String:” prompt will never exceed 30 characters, even if the actual input string can be larger. The scrolling step is currently 10 characters.
- If the entered filename length exceeds 30 characters, these fields will only display the final 30 characters of the filenames entered, indicating that it is truncated by displaying three dots before the actual filename (...name).

7.2 Switch Clustering Enhancements

Switch Clustering using Split MultiLink Trunking (SMLT) and Single Link Trunking (SLT) works in standalone and stacked environments in Release 5.1. Nortel supports the following configurations in Release 5.1:

- Triangle – both standalone and stack
- Square – both standalone and stack

Stack Switch Clustering allows up to a full stack of eight switches to have interswitch trunking (IST) connections to another stack of switches, which provides greater redundancy and bandwidth aggregation. Square configuration allows additional, more complex and resilient network designs. Nortel supports two square configurations:

- Standalone 5500 Series IST peers connected to a second set of standalone 5500 Series IST peers
- Standalone 5500 Series IST peers connected to 8600 IST peers
- Stacked 5500 Series IST peers connected to 8600 IST peers (Layer 2 on the 5500 only with Layer 3 on the 8600)

Connection from the “core” of a square configuration requires static routes pointing to the virtual router (VR) IP address.

Additionally, due to a known issue in v5.1 software, devices SLT connected to the 5500 IST peers should have autonegotiation disabled and hard coded to a fixed speed, or configure Customized Auto-Negotiation Advertisements (CANA) on the 5500 IST peers.

7.3 Problems Resolved in This Release

When SNMPv3 is configured and the base unit rebooted, the SNMPv3 view is lost (**Q01843759**).

After a reboot, IP routing was enabled by itself on management VLAN (**Q01800968).**

Under certain conditions, the number of MAC addresses displayed was not correct (**Q01793281-01**).

SMLT did not transition to NORMAL state when VLACP went down on links (**Q01798409**).

In a certain setups, a MAC address was incorrectly displayed in two VLANs (**Q01812965**).

With 55xx used as an aggregation switch, sometimes the traffic was not properly forwarded to the edge (**Q01823398**).

AUR was not ready with temporary base unit (**Q01835087**).

Large amount of login timeout logs was recorded (**Q01838853-01**).

When primary radius server was not available, the EAPoL re-authentication failed (**Q01787801**).

With autosave disabled and "copy config nvram" command issued quickly after the last configuration command, the latest DHCP changes did not get saved in the configuration **(Q01837653-01).**

With a specific SSH client, an SSH session to the switch could not be established (**Q01837389-04**).

ARP spoofing does not protect against a gratuitous ARP from a device configured with the gateway IP address. (**Q01849206-01**).

STP status for MLT ports was unexpectedly changed after reboot (**Q01808492**).

In a two-unit stack, a loss of routing could have resulted if a unit was powered off **(Q01606521-01).**

A NEAP device was not able to communicate through ERS5500 while the device was authenticating (**Q01833016-03**).

OSPF adjacencies failed when 'Unknown Multicast Filter' was enabled (**Q01828889**).

Under certain circumstances, a PC with NEAP Radius Authentication could lose network connectivity although an apparently authenticated port (**Q01784034**).

Telnet connectivity not available for a short period of time after the IP address of the switch was changed (**Q01812713-02**).

When using MHMA mode, the RADIUS VLAN attribute was ignored after reset (**Q01837436**).

IGMP General Query was not sent after reboot (**Q01491178-04**).

First EAP ID response was dropped (**Q01769619-01**).

After reboot, ports move from IGMP configured VLAN to non-IGMP VLAN (**Q01778281-01**).

Specific DHCP request packet could not be forwarded (**Q01857659**).

IPFix template updates were not sent to the Report Analyzer (**Q01760099-01**).

Specific DHCP discover packet caused an exception (**Q01878953**).

DHCP Snooping Filters could not be re-enabled (**Q01851019**).

When a switch in a stack rejoined after a reboot, some or all FTP sessions on non-rebooted units could be dropped (**Q01720501-03**).

When the base unit of a 2-unit stack rebooted, the LACP link on a non-base unit stopped forwarding the traffic **(Q01733917-02).**

Under certain conditions, TFTP transfers caused stack corruptions (**Q01881540**)

ARP entries were not learned when receiving DHCP Discover packets from a device that had multiple VLANs using the same MAC address (**Q01768780-01**).

Erroneous messages were recorded in the system log if DHCP Snooping was enabled (**Q01835111-01**).

File sharing among Wireless users failed when the DHCP snooping was enabled (**Q01859724-01**).

Under certain conditions, SSH authentication failed with passwords of 16 or more characters (**Q01900077**).

After several authentication tries, a PEAP client failed to be authenticated (**Q01862479-01**).

NEAP clients could not communicate when Radius assigned VLAN option was configured (**Q01877389-01**).

Erroneous fan failure reports, while the fans were operational (**Q01586344-01**).

The switch routed traffic to the default route while a specific route was in its routing table (**Q01898900**).

Stack intermittently crashed after a soft reset of one unit during the AUR process (**Q01852682-01**).

VLACP ports were flapping when the DHCP snooping and the ARP inspection are enabled (**Q01860016**).

For a particular NSNA setup, some switch ports may become disabled (**Q01927868**).

8. Outstanding Issues

.
In a Stack of 3 with STP operation mode of RSTP; Pings over SMLT dropped during the reset of unit 2 (Q01839508).

OSPF multicast addresses should not be configured for unknown-mcast-allow-flood with OSPF enabled (Q01900472).

IGMP general queries were not sent over MLT after a reboot (Q01900952).

When disabling IGMP and then re-enabling it, the General queries were not sent over MLT (Q01900995)

ARP Spoofing did not work on MLT/LACP (Q01914799)

9. Known Limitations

In order to prevent certain users from advertising their own MAC address as the MAC address of the default gateway, the execution order for QoS ARP Spoofing policies was changed by giving the highest priority to the actual 3rd policy (as described in NN47200-504, "Configuration - Quality of Service" - page 41), the one that drops all ARP packets with a source IP address equal to the identified default gateway. The new execution order for QoS ARP Spoofing policies will be:

1. Drop all ARP packets with a source IP address equal to the identified default gateway.
2. Pass all broadcast ARP requests.
3. Drop all non-broadcast ARP requests.
4. Drop all ARP packets with a target IP address equal to the identified default gateway.
5. Pass all ARP responses

10. Documentation Corrections

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 .

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