



# Ethernet Routing Switch 5510/5520/5530 Software Release 5.0.8

## 1. Release Summary

Release Date: 14-February-2008

Purpose: Software patch release to address customer requirements and customer found 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 5.0" (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_500003_diag.bin | Diagnostic image       | 811,864           |
| 5530_508034.img      | Agent code image       | 5,410,536         |
| 5530_508035s.img     | Agent code image (SSH) | 5,636,108         |

## 5. Version of Previous Release

Software Version 5.0.7

## 6. Compatibility

This software release is managed with Java Device Manager (JDM) release 6.0.2.

## 7. Changes in This Release

### 7.1. New Features in This Release

#### 7.1.1. Port-mirroring on 802.1x (EAP) port

Due to potential security risks, an EAP port could not be monitored or a mirror port, and a mirror port could not be set as an EAP port.

An enhancement is provided in this release to enable port-mirroring on EAP ports. Three new global commands to enable or disable port mirroring on EAP ports are provided under the EAP sub-commands/group. The format of the commands is:

```
EAPo1 allow-port-mirroring  
no EAPo1 allow-port-mirroring  
default EAPo1 allow-port-mirroring
```

The user will be provided with a warning message about the potential security risk in each one of the following situation:

- When enabling the feature
- If the user is mirroring an EAP port or try to enable EAP on a mirror port
- When disabling the feature and there are still mirroring ports with EAP enabled

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

In this release, radius assigned VLAN will be allowed upon non-eap radius-mac-authentication, in MHMA mode. Note that radius-mac-authentication feature must be enabled, for getting any radius-assigned values for non-eap macs.

When this new feature is not enabled, VLAN attributes from radius server will be ignored.

For this feature to be enabled, it has to be enabled both globally and at the interface (port).

When this feature is enabled, behavior will be exactly as for a radius assigned VLAN for EAP authorized MACs, except that in this case, it would be for radius authorization for a non-EAP MAC.

#### CLI commands:

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

```
EAPo1 multihost use-radius-assigned-VLAN  
no EAPo1 multihost use-radius-assigned-VLAN  
default EAPo1 multihost use-radius-assigned-VLAN  
  
EAPo1 multihost non-EAP-use-radius-assigned-VLAN  
no EAPo1 multihost non-EAP-use-radius-assigned-VLAN  
default EAPo1 multihost non-EAP-use-radius-assigned-VLAN
```

### 7.1.3. RADIUS request for ADAC MACs

ADAC MACs are allowed on an EAP enabled port subject to the following requirements:

Global EAP is enabled and Interface EAP is enabled (status = auto)

Global EAP multihost for configured NonEAP MACs is allowed

Interface EAP multihost for configured NonEAP MACs is allowed

ADAC is globally enabled

ADAC is locally enabled as a Telephony Port

Uplink and CallServer Ports cannot be non-EAP-enabled at the same time.

Similarly, ADAC cannot be enabled locally on a port that is EAP enabled but is not configured for multi-host and to allow non-EAP MACs.

When a new MAC is seen on an EAP port and is allowed based on ADAC credentials, a dummy radius request as for a non-EAP radius MAC authentication will be sent. This is for servers to learn the MAC to update their databases. Any response from radius for an ADAC authenticated MAC will be ignored.

### 7.1.4. 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 the CPU load on the specific unit. The values show the cpu utilization for the last 10 seconds, 1, 10 and 60 minutes, for the last 24hours and since system bootup.

Normal values for the cpu utilization are less than 30% cpu load. In a stack configuration this value is around 30%, up to 35% cpu load on higher stack configurations. 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 boot the memory utilization should be around 99%; when the switch performs some operations after the boot, for example download, the available memory value will decrease; after the operations end, there will be an increase to the initial value in the available memory value; the low mark will save the lowest available memory since bootup. In normal switch operation there should always be available memory for applications to use.

**NOTE:** The CPU utilization may vary due to different configurations at the customer networks. It is advisable 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. The customers should debug their switches/networks using the other available methods.

### 7.1.5. 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, if 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.

## **7.1.6. 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*

| <i>Unit #</i> | <i>Last Configuration-Save Time-Stamp</i> | <i>Ready For Replacement</i> |
|---------------|-------------------------------------------|------------------------------|
| <i>1</i>      | <i>3 days 10:23:02</i>                    | <i>Yes</i>                   |
| <i>2</i>      | <i>0 days 00:01:40</i>                    | <i>No</i>                    |
| <i>3</i>      | <i>3 days 10:12:33</i>                    | <i>Yes</i>                   |
| <i>6</i>      | <i>3 days 10:12:34</i>                    | <i>No</i>                    |
| <i>8</i>      | <i>3 days 10:12:35</i>                    | <i>Yes</i>                   |

## **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.

## **7.2. Old Features Removed From This Release**

None.

## **7.3. Problems Resolved in This Release**

**Q01784387** - Improved the DMLT recovery algorithm per customer request

**Q01778281** - ERS: 5520 - Ports move from IGMP configured VLAN to non-IGMP VLAN

**Q01759628** - ERS 5520: MLT Uplink ports are not recovering after power cycle of base unit

**Q01655126** - 55xx: Request that IP Manager ACLs support SSH

**Q01798679** - ERS 5510/5520/5530 DHCP 5.0.x code DHCP snooping causes connectivity issues

## **8. QoS Queue Assignments**

In order to optimize certain types of traffic flow, some changes to QoS queue assignments were made in 5.0.2 release. The changes are as follows:

```
qos queue-set-assignment queue-set 2 1p 4 queue 2
qos queue-set-assignment queue-set 2 1p 5 queue 2
qos queue-set-assignment queue-set 3 1p 4 queue 3
qos queue-set-assignment queue-set 3 1p 5 queue 3
qos queue-set-assignment queue-set 4 1p 2 queue 4
qos queue-set-assignment queue-set 4 1p 3 queue 4
qos queue-set-assignment queue-set 4 1p 4 queue 4
qos queue-set-assignment queue-set 4 1p 5 queue 3
qos queue-set-assignment queue-set 4 1p 7 queue 2
qos queue-set-assignment queue-set 5 1p 2 queue 5
qos queue-set-assignment queue-set 5 1p 5 queue 5
qos queue-set-assignment queue-set 5 1p 6 queue 1
qos queue-set-assignment queue-set 5 1p 7 queue 2
qos queue-set-assignment queue-set 6 1p 3 queue 4
qos queue-set-assignment queue-set 6 1p 4 queue 6
qos queue-set-assignment queue-set 6 1p 6 queue 1
qos queue-set-assignment queue-set 6 1p 7 queue 2
qos queue-set-assignment queue-set 7 1p 6 queue 1
qos queue-set-assignment queue-set 7 1p 7 queue 2
qos queue-set-assignment queue-set 8 1p 6 queue 1
qos queue-set-assignment queue-set 8 1p 7 queue 2
```

## 9. **Known Limitations**

### **EAP:**

1. When using radius assigned VLANs, as a general rule, set VLAN configcontrol to flexible (not the default strict which restricts it to 1 VLAN)
2. When EAP ports are mirrored, a potential security risk may be created. Please take necessary actions to assure data protection before enabling port mirroring on EAP ports.
3. When ADAC and EAP are enabled on the same port and an ADAC authenticated MAC is seen on the port, a radius request is sent to the radius server. Even if Radius rejects the MAC address, the MAC will be allowed if already ADAC-authenticated. The expectation is radius will ignore or reject the request, since it is an ADAC MAC.
4. When using a guest VLAN, it is required for the switch/stack to have an IP address. Otherwise the ports will not be moved to the guest VLAN.

### **DMLT (CR Q01784387):**

1. It is required that all members of a DMLT group have the same configuration parameters – VLAN, VLACP configuration, IGMP and rate-limiting. When the stack is formed, the MLT application performs a consistency check on all members of a MLT/DMLT group for mismatched configuration parameters. Prior to release 5.0.8, if any DMLT member port is found with an inconsistent configuration, all members of that DMLT will be disabled.
2. In release 5.0.8 and later, the MLT application will shutdown only the port which has the different configuration. In some cases, the code is unable to determine which port is incorrectly configured. The MLT application will keep the DMLT port on the base unit enabled, and disable DMLT ports on non-base units. Leaving this single link active assures that connectivity to the switch/stack is not lost.

### **AUR Enhancement:**

1. Before disabling the AUR auto-save feature, the user should create all the VLANs that are required in the stack. If the user disables AUR and creates VLANs in a stack and then restores the old configuration to the NBU, the port members of all the VLANs except the default VLAN will be removed from that unit.
2. Before disabling the AUR auto-save feature, the user should make sure all the configurations related to “ARP inspection” are already made. If the user disables the AUR auto-save and then modifies the ARP inspection on a port, those changes will not be reflected if the unit on which the configurations are modified is replaced.
3. If the base unit is reset before user attempts to restore the saved configuration, user will lose all the saved configuration of the non-base units

## 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> .

---

Copyright © 2008 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, and the Globemark 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>