Part No. 216019-A
December 2003

4655 Great America Parkway
Santa Clara, CA 95054

Release Notes for the BayStack 380-24T and BayStack 380-24F Switch, Software Version 3.0

216019-A

NORTEL NETWORKS™

Copyright © 2003 Nortel Networks

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 Networks Inc.

Trademarks

Nortel Networks, the Nortel Networks logo, the Globemark, Unified Networks, BayStack 380, and Optivity are trademarks of Nortel Networks.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

Java is a trademark of Sun Microsystems, Inc.

Macintosh is a trademark of Apple Computer, Inc.

Netscape Navigator is a trademark of Netscape Communications Corporation.

Restricted rights legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Contents

Introduction

These release notes document the known issues and new features of BayStack 380, software release 3.0 for the BayStack 380-24T 10/100/1000 switch and the BayStack 380-24F Gigabit switch.

New Features for BayStack 380, software release 3.0

The new features for BayStack 380, software release 3.0 include:

- BayStack 380-24T and BayStack 380-24F common image
- BayStack 380 command line interface (CLI)
- ASCII configuration download
- 10/100 Mbps management port
- Class of service
- Virtual Area Networks (VLANs)
- IGMP snooping
- Single fiber fault detection
- ASCI configuration generator

Hardware compatibility

BayStack 380 Software Version 3.0 is compatible with the following Nortel Networks products:

- BayStack 380-24T 10/100/1000 Switch
- BayStack 380-24F Gigabit Switch

BayStack 380 sofware version 3.0 compatibility matrix

The components for the BayStack 380 Software Version 3.0 are:

- BayStack 380 Runtime Image Software Version 3.0.0.x (bs380_boss30039.img)
- BayStack 380 Boot / Diagnostic Software Version 3.0.0.x (bs380_boss3002_diag.bin)
- Java Device Manager software version 5.7.4.0 (jdm_5740)
- BayStack 380 Management Information Base (MIB) definition files (bs380mibs_3020xxx.zip)

BayStack 380-24F Software Version 2.1 Compatibility Matrix

The components for the BayStack 380-24F, Software Version 2.1 are:

- BayStack 380 Runtime Image Software Version 2.1.0.x (bs380_210x.img) and Software Version 2.1.1.x (bs380_211x.img)
- BayStack 380 Boot / Diagnostic Software Version 2.1.0.x (bs380diags_210x.bin)
- Java Device Manager version 5.6.1.0 (jdm_5610)
- BayStack 380 Management Information Base (MIB) definition files (bs380mibs_100xxx.zip)

BayStack 380-24T Software Version 2.0 Compatibility Matrix

The components for the BayStack 380 Software Version 2.0 are:

- BayStack 380 Runtime Image Software Version 2.0.0.x (bs380_200x.img) and Software Version 2.0.1x (bs380_201x.img)
- BayStack 380 Boot / Diagnostic Software Version 2.0.0.x (bs380diags_200x.bin)

- Java Device Manager version 5.5.6.0 (jdm_5560)
- BayStack 380 Management Information Base (MIB) definition files (bs380mibs_100xxx.zip)

Known issues for BayStack 380, software version 3.0

BayStack 380, software version 3.0 has the following known issues:

- When an SFP GBIC is inserted into port 24 of the BayStack 380-24T, the corresponding copper port (port 24) may drop link even if there is no link on the GBIC port. Therefore, only insert a GBIC into port 24 if the corresponding copper port is not in use (Q00646868).
- When using Device Manager to add ports to VLANs, create the VLAN first, and then create and add the ports to the VLAN. If you try to create the VLAN and add the ports to the VLAN in one step, the PVID value will not reflect the value of the VLAN if auto-PVID is enabled. (Q00322230)
- You should have autonegotiation enabled or disabled on both ends when connecting a BayStack 380 switch to another device. (Q00732308)
- Before making changes to MAC address security settings, MAC address security should be disabled, changes should be made, and then re-enabled. (Q00554340)
- When using the ASCII Configuration Generator, the SNMP community names are not retained. This prevents the security of the switch from being compromised. (Q00739786)
- Unknown packets may be counted incorrectly as flooded packets on a port that is blocked by spanning tree. (Q00679795)
- In the Port Mirroring section of *Using the Baystack 380 10/100/1000 Switch* (212859-A), Table 30 lists Monitoring Mode as an option. This option is not available. (Q00689562)
- When generating an ASCII Configuration File using the ACG function via CLI, uploading the ASCII Configuration File must be done using CLI. (Q00725241)
- When using console, CLI, JDM, WEB and SNMP, all ports can be removed from the management VLAN with no warning about losing connectivity. (Q00688737)

- If all 512 VLANs are configured and all ports are in all VLANs, ping responses may take up to 50 seconds during the configuration change. (Q00733624)
- Copy running-config fails when using tftp server from OSM 3.0 (Windows) when VLANs are configured (Q00733629)
- The maximum number of characters for RADIUS Shared Secret field that you can set using console or telnet interface is 15 characters.
- The maximum number of characters for RADIUS Shared Secret field that you can set using Web interface is 16 characters. (Q00598639)
- On BayStack 380-24F, the status of the negotiated flow control is displayed incorrectly, but it does operate correctly.(Q00803489)
- The BayStack 380-24T is not recommended for use in environments requiring low latency transmissions (e.g. VoIP). (Q00764281)

Issues corrected in BayStack 380, software version 3.0

The BS380-24T lost IP management traffic. (Q00705556)

Packets larger than 1521 bytes are improperly treated. (Q00651076)

Topology packets are only sent out of the primary MLT port. (Q00623080)

New Features for BayStack 380, software version 3.0

The following sections document the new features for BayStack 380, software version 3.0

BayStack 380-24T and BayStack 380-24F common image

The BayStack 380-24T and BayStack 380-24F common image provides customers with a single software image for both platforms: BayStack 380-24T (Copper BayStack 380) and 380-24F (Fiber BayStack 380). This common image provides the same software to all user interfaces such as menu, telnet, CLI, WEB, JDM.

When the system boots up, software version 3.0 (or higher) detects the 380 platform type and runs accordingly. This is transparent to the user and no additional configuration is required.

BayStack 380 command line interface (CLI)

BayStack 380, software release 3.0 includes a command line interface (CLI). This interface is documented in the *Reference for the BayStack 380-24T and 24F Command Line Interface, software version 3.0*.

ASCII configuration download

This feature downloads an ASCII configuration file from a TFTP server, parses the commands in the file, and configures the switch based on the commands in the file.

You can edit the configuration file on any host computer using a regular text editor. This provides quick configuration of the switch, using an easy-to-modify text configuration file.

You use the Nortel Networks Command Line Interface (NNCLI) to make changes to the ASCII configuration file.

ASCII Configuration File Download screen

The ASCII Configuration File Download screen (Figure 1) allows you to specify the IP address of an TFTP server and the name of the configuration file on the server, and to download an ASCII configuration file containing CLI commands from the TFTP server to configure the switch.

1 Choose ASCII Configuration File Download from the Configuration File.

The Menu opens to the ASCII Configuration File Download screen.

Figure 1 ASCII Configuration File Download screen

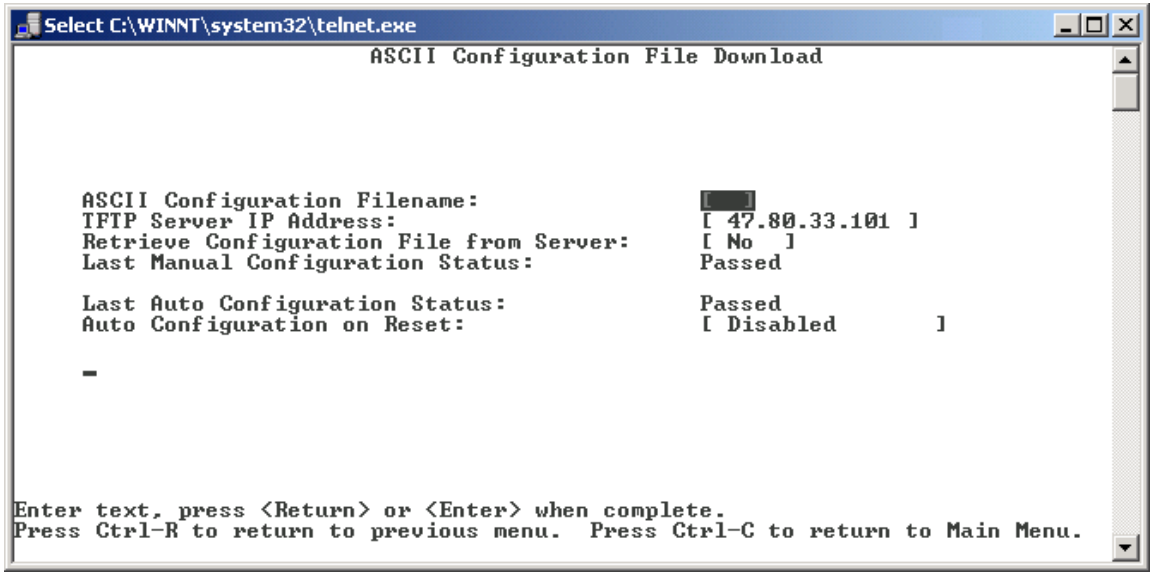

Table 1 describes the ASCII Configuration File Download screen fields.

Table 1 ASCII Configuration File Download screen fields

Field	Description
ASCII Configuration Filename	Enter the file name you have chosen for the ASCII configuration file. Choose a meaningful file name that will allow you to identify the file for retrieval when required. The file must already exist on your TFTP server and must be read enabled. Default Value Zero-length string Range An ASCII string of up to 30 printable characters
TFTP Server IP Address	The IP address of your TFTP load host. Default Value 0.0.0.0 (no IP address assigned) Range Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Retrieve Configuration File from Server	Specifies whether to retrieve the stored switch ASCII configuration file from the specified TFTP server (default is No). If you choose Yes, the download process begins immediately and, when completed, causes the switch to be configured according to the CLI commands in the file. Use the spacebar to toggle the selection to Yes. Press [Enter] to initiate the process. Default Value No Range Yes, No
Last Manual Configuration Status	The system displays if the last manual configuration passed or failed. Default Value Passed Range Passed, Failed
Last Auto Configuration Status	The system displays if the last automatic configuration passed or failed. Default Value Passed Range Passed, Failed
Auto Configuration on Reset	Allows you to choose to Disabled, Use Configured, or Use BootP: • Disabled—Auto configuration on reset is disabled. • Use Configured—Use manually configured ASCII configuration filename and TFTP server address for auto configuration on reset. • Use BootP—Retrieve ASCII configuration filename, and optionally server address, using BootP, when BootP is enabled, and perform auto configuration on reset using these parameters. Default Value Disabled Range Disabled, Use Configured, Use BootP

The ASCII configuration file does not generate any output to the Console. If there is an error in the configuration file, the NNCLI parser exits on the first error found and displays an error message.

When you select the 'Yes' option in the 'Retrieve Configuration File from Server' menu item, the file is transferred in ASCII format from the server, via TFTP, and stored in local memory.

There are three Auto-Configuration modes:

- Disabled (default)
- Use Configured
- Use BootP

The 'Use Configured ' mode uses the File Name and TFTP Server Address parameters in the “ASCII Configuration File Download screen” to retrieve an ASCII configuration File automatically after reset, and configure the switch or stac.k

When the switch is powered up, there is a delay while the system initializes, resolves its IP address, spanning tree states go into forwarding, before the switch attempts to retrieve the configuration file. Appropriate messages are displayed during initialization.

The 'Use BootP' mode retrieves the ASCII Configuration file name, and the TFTP server IP address if the file is not on the BootP server, from the BootP response, when the switch performs BootP on reset for its IP configuration.

If this mode is selected, but no ASCII configuration file is specified in the bootptab file using the bf keyword, a message "no boot file found" is printed by the kernel on the base unit.

The configuration file server address can be specified in the bootptab file using the cs keyword. If the file server address is not specified in the bootptab file, but a valid file name is specified, it is assumed that the file is on the BootP server. If the file is not found, or not accessible, the TFTP process times out and a message "Transfer Timed Out" is printed on the base unit screen by the kernel.

BootP is a basic bootstrap protocol implemented over the Internet User Datagram Protocol (UDP). It allows a booting target to configure itself dynamically by obtaining its IP address, subnet mask, gateway IP address, boot file name and the boot hosts IP address over the network. It thus allows centralized management of the target boot parameters on the host system.

After the ASCII configuration file name is retrieved using BootP, the file is transferred in ASCII format from the server, via TFTP, and stored in local memory. The NNCLI parser parses the commands in the file and configures the switch.

10/100 Mbps management port

The 10/100 management port feature for the BayStack 380-24F (fiber) switch allows the switch to be managed using the RJ 45 port without the need for a fiber connection to the switch.

Each packet whose destination MAC address matches the switch mac address or the broadcast mac address is steered to CPU and processed by the IP stack of the underlying OS. This enables the possibility to establish an IP communication between external hosts and the BayStack 380-24F unit.

Communication with an external device is allowed as long as the remote end's MAC address is not learned on the front panel ports. If such a case is detected then the ingress frames whose mac addresses are found in the switch's MAC address forwarding table are dropped – a situation like this indicates that the remote end is connected to one of the front panel ports (loop) and the IP communication will be done using that port.

The 10/100 management port does not participate in the STP protocol and is isolated from the front panel ports – there is no a way to switch traffic directly between the RJ 45 management port and the SFP GBIC ports.

The management port is set to auto negotiate the speed and duplex capabilities. It may operate at 10 or 100 Mbps and full or half duplex. You cannot configure the speed and duplex setting of the management port.

To protect the CPU from excessive traffic and broadcast storms, the management port is limited to receiving 200 packets per second. If the receive rate on the management port should exceed this threshold, the receive circuitry on the port will be disabled for 1 second and then re-enabled.

Class of Service

BayStack 380, software version 3.0 provides Class of Service DSCP priority mapping in addition to 802.1p priority mapping. Class of Service (COS) is the process of forwarding frames based on their priority derived either from a packet header (VLAN tag, DS field within the IP header) or a programmable field within the bridge entity.

The BayStack 380 provides the capability to map the DSCP values in the IPv4 frame to any of four (4) priority queues. If the frame is not an IPv4 frame, DSCP mapping is not applicable and the 802.1p mapping is used. The BayStack 380 allows the user to map all 64 possible DSCP values to any of the four priority queues. The IPv4 frames will be forwarded according to the COS queue priorities.

The CoS/DSCP mapping feature may be globally enabled or disabled through the Command Line Interface (CLI) and the Console interface. The COS scheduler is only supported through the CLI and the Console Interface.

This feature does not support the experimental bits (two highest order bits) of the Diffserv byte.

The COS scheduler can be configured to work in two modes:

- Strict Round-Robin

  Frames in the higher COS queues are transmitted first

- Weighted Round-Robin

  Frames in all the queues are transmitted in a Round Robin fashion according to the weights values assigned to each COS queue

You can set the DSCP to CoSq from the console screen as follows.

Figure 2 Class of Service

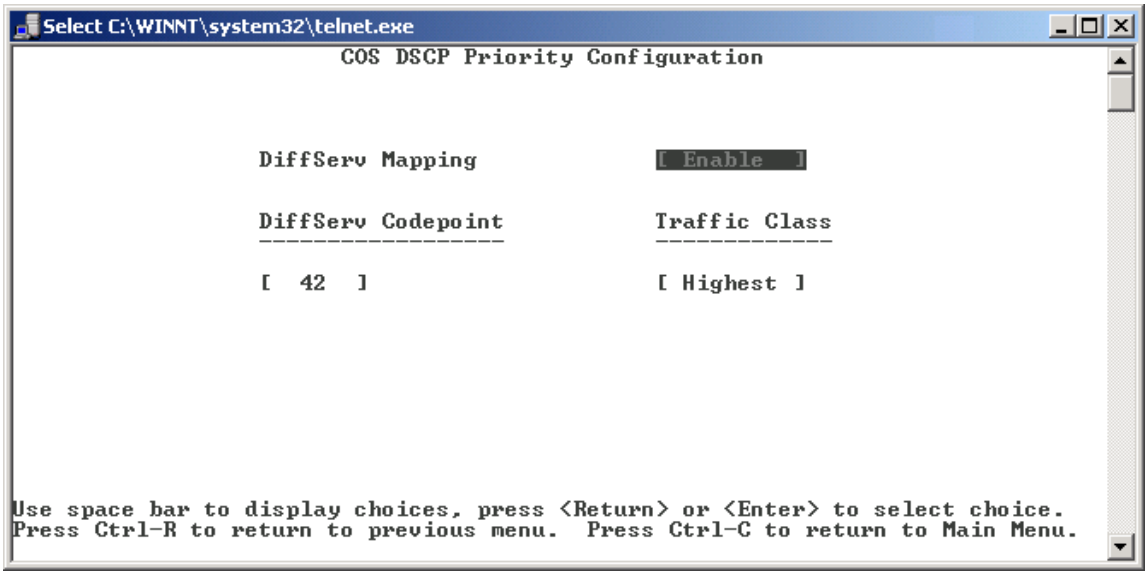

The COS DSCP screen is under the 'Traffic Class Configuration' menu screen.

Table 2 Class of service fields

Field	Description
Diffserv Mapping	Enables or disables the COS/DSCP mapping. Values Enable or Disable
Diffserv Codepoint	Specifies the codepoint value to be mapped to a specific queue. Values 0 to 63
Traffic Class	Specifies the COS queue. Values Low, Medium, High, Highest

By default:

- All Diffserv Codepoints are mapped to the "Low" CoS queue
- The default for mappinng of Diffserv Codepoints to the CoS queues is disabled.

512 Virtual Local Area Networks (VLANs)

VLANs allow multiple broadcast domains on a device. This can decrease the amount of traffic going through the device, which means fewer collisions and better network performance. Encapsulating certain users and devices on a VLAN allows you to design a more efficient network.

BayStack 380 has a tagging option from 802.1Q that allows users to specify which ports are to be tagged or not. This allows the device to add tags to untagged traffic on tagged ports and strip tags off of tagged traffic on untagged ports.

The device is part of one VLAN called the management VLAN. Only traffic on this VLAN is seen by the device, such as ping or SNMP. Traffic on other VLANs are not be able to communicate with BayStack 380.

You can create and configure up to 512 port-based VLANs on BayStack 380. BayStack 380 supports up to 4000 VLAN entries, but only a maximum of 512 can be active at any given time. BayStack 380 also supports only Independent VLAN Learning (IVL) VLANs, it does not support Shared VLAN Learning (SVL) VLANs.

You cannot delete the default VLAN, and there is only one Management VLAN and any given time. The Management VLAN can be changed and the current Management VLAN cannot be deleted.

The switch is a member of the Management VLAN. PVIDs can be assigned on a per port basis. There is also an AutoPVID option that automatically assigns a PVID to an untagged port when the port is added to a VLAN.

Tagging is available on a per port basis. You can configure tagged and untagged ports on the same VLAN.

IGMP snooping

On a single physical network, multicasting communication is very simple. The sending process specifies a destination (multicast) address, which gets converted into the corresponding Ethernet address and the datagram appears on the cable. The receiving processes (usually there are multiple receivers for a given multicast address) - which have notified the underlying IP layers that they want to receive datagrams for a given multicast address - receive a copy of this datagram.

When multicasting is used on more than one physical network and the multicast datagrams have to pass through routers, there is a need for a special protocol - the Internet Group Management Protocol (IGMP). IGMP is used by IP hosts to report their multicast group memberships to any immediately-neighboring multicast routers. IGMP is part of the IP layer, and is required to be implemented by all hosts wishing to receive IP multicasts. IGMP messages are encapsulated in IP datagrams, with an IP protocol number of 2.

IGMP snooping enables the switch to selectively forward multicast traffic only onto ports where particular streams are expected.

- Pruning multicast streams from non-participating ports improves network performance.
- IGMP V1 and V2 messages (Queries, Reports, Leaves) are intercepted and processed
- A database of learned “IGMP speaking” ports is maintained per VLAN

Individual multicast streams are sent only to ports which have received reports for the corresponding multicast groups. Proxy functionality permits consolidating multiple reports received for the same group into a single report and sending the report to any existing router port.

The switch can interoperate with both IGMP V1 and V2 hosts/routers. Snooping is performed on the IGMP messages passed around through the network. Pruning prevents multicast streams from going out to ports that are not interested in them. IGMP port aging terminates stale connections, avoiding sending streams to hosts which do not listen for them anymore. IGMP Snooping and Proxy services are enabled in the default configuration.

Parameters which can be adjusted per VLAN:

- Snooping (Enabled/Disabled – default Enabled) – switch; other settings depend on this to be enabled
- Proxy (Enabled/Disabled – default Enabled) – toggles proxy functionality. When ‘off’, all incoming reports are forwarded to router ports
- Robust Value (0-64 – default 2) – how many times a port is queried for group participation before aging it out
- Query Time (1-512 – default 125) – number of seconds between each query upon port aging.
- Static Router Ports V1/V2 – ports manually set as router ports (only useful when used in an environment with old equipment which does not send queries)

A Robust Value of ‘0’ is a convention used to activate ‘Fast Leave’ mode. The default behavior is to query the leaving port for ‘Robust Value’ times before aging it out. With ‘Fast Leave’, the leaving port is dropped immediately from the group membership. The user interface permits changing all the available IGMP parameters. The ‘Set Router Ports’ field is used for specifying the IGMP version of the Static Router Ports.

A command-line interface is also available for IGMP. Commands are accessed with the ‘vlan igmp’ prefix.

IGMP produces three types of messages, which are sent between IGMP routers and hosts:

1. Membership Query messages - sent periodically by querier routers in order to:
   - determine which multicast groups have members learn the groups that have members on an attached network. (General Query)
   - learn if a particular multicast group has any members on an attached network (Group-Specific Query).

   Membership Query messages are usually referred to as "Query" messages, or "queries".
2. Membership Report messages - sent by IGMP hosts as response to IGMP queries, or when they first attach to a multicast group. Also referred to as "reports".
3. Leave Group messages - or simply "leave" messages. Introduced by IGMPv2, they allow group membership termination to be quickly reported to the routing protocol.

In a bridged Ethernet environment, IP multicast is directly mapped to broadcast transmissions. Every IP multicast packet is forwarded on all links of a layer 2 device, such as an Ethernet switch, and delivered to all segments of an extended LAN. Network performance degrades as a network carries more broadcast traffic. End stations, which are not interested in particular IP multicast streams, are offered the same load as the rest of the network indiscriminately.

The increasing application of IP multicast has lead to the development of new techniques for optimizing multicast by layer 2 devices. IGMP snooping is such a technique whereby a layer 2 switch selectively forwards multicast traffic only onto ports where particular IP multicast streams are expected. A switch can identify those ports by "snooping" for IGMP communication between routers and hosts.

The BayStack 380 hardware can be configured to intercept all IGMP traffic and forward the corresponding packets to the CPU. The IGMP snooping application will be responsible for setting forwarding/filtering masks for the front panel ports with respect to the learned multicast addresses.

IP multicast traffic will only be forwarded to those ports from which IGMP queries and reports have been received; multicast packets will be dropped for the rest of the network ("pruning"). As a result, network performance is improved. It should be pointed out that a BayStack 380 switch is neither an IGMP router nor an IGMP host. It is a layer 2 device located between routers and hosts, which is able to prune and optimize IP multicast in an Ethernet environment.

IGMP snooping and Multicast

BayStack 380 Switches can sense Internet Group Management Protocol (IGMP) host membership reports from attached stations and use this information to set up a dedicated path between the requesting station and a local IP Multicast router. After the pathway is established, the BayStack 380 Switch blocks the IP Multicast stream from exiting any other port that does not connect to another host member, thus conserving bandwidth. The following section describes how BayStack 380 Switches provide the same benefit as IP Multicast routers, but in the local area.

IGMP is used by IP Multicast routers to learn about the existence of host group members on their directly attached subnets (see RFC 2236). The IP Multicast routers get this information by broadcasting IGMP queries and listening for IP hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IP Multicast source that provides the data streams and the clients that want to receive the data.

Figure 3 shows how IGMP is used to set up the path between the client and server. As shown in this example, the IGMP host provides an IP Multicast stream to designated routers that forward the IP Multicast stream on their local network only if there is a recipient.

Figure 3 IP Multicast propagation with IGMP routing

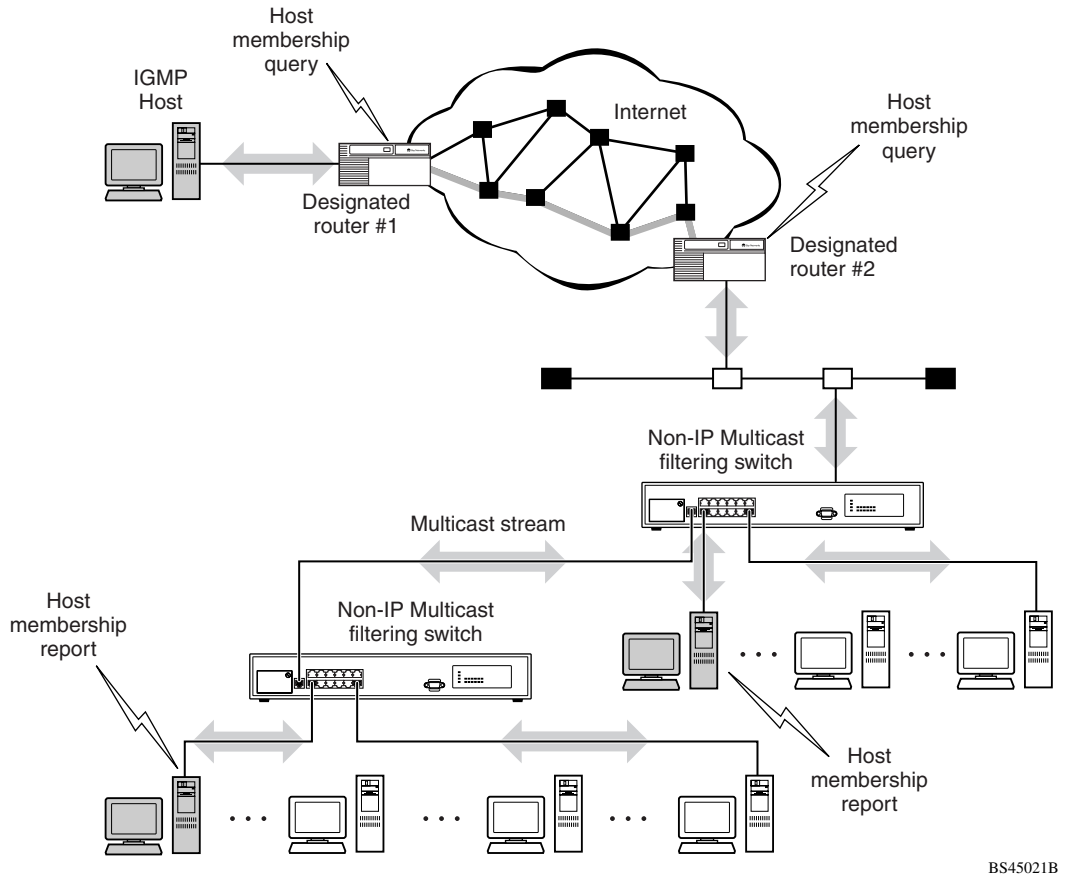

The client/server path is set up as follows:

1. The designated router sends out a host membership query to the subnet and receives host membership reports from end stations on the subnet.
2. The designated routers then set up a path between the IP Multicast stream source and the end stations.
3. Periodically, the router continues to query end stations on whether or not to continue participation.
4. As long as any client continues to participate, all clients, including nonparticipating end stations on that subnet, receive the IP Multicast stream.

Note: Although the nonparticipating end stations can filter the IP Multicast traffic, the IP Multicast traffic still exists on the subnet and consumes bandwidth.

IP Multicast can be optimized in a LAN by using IP Multicast filtering switches, such as the BayStack 380 Switch.

As shown in Figure 3, a non-IP Multicast filtering switch causes IP Multicast traffic to be sent to all segments on the local subnet.

The BayStack 380 Switch can automatically set up IP Multicast filters so the IP Multicast traffic is only directed to the participating end nodes (see Figure 4).

Figure 4 BayStack 380 Switch filtering IP multicast streams (1 of 2)

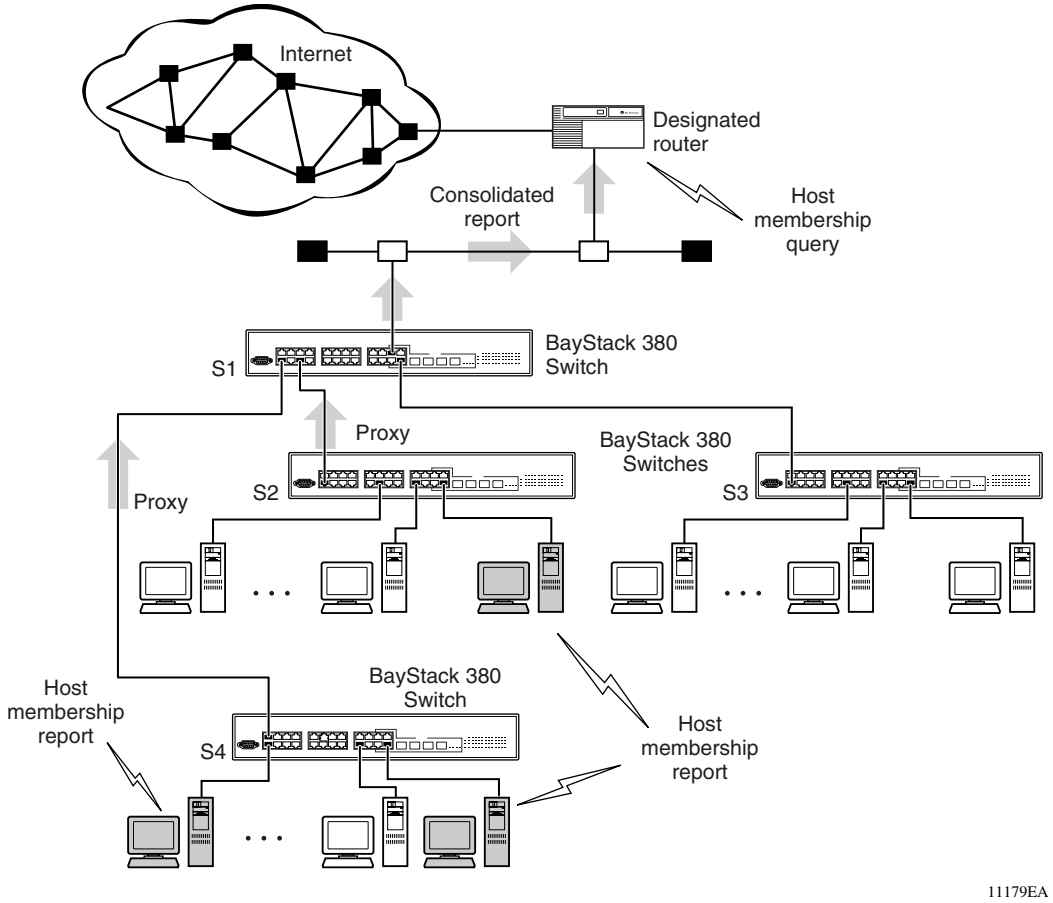

In Figure 4, switches S1 to S4 represent a LAN connected to an IP Multicast router. The router periodically sends Host Membership Queries to the LAN and listens for a response from end stations. All of the clients connected to switches S1 to S4 are aware of the queries from the router.

One client, connected to S2, responds with a host membership report. Switch S2 intercepts the report from that port, and generates a *proxy* report to its upstream neighbor, S1. Also, two clients connected to S4 respond with host membership reports, causing S4 to intercept the reports and to generate a *consolidated proxy report* to its upstream neighbor, S1.

Switch S1 treats the consolidated proxy reports from S2 and S4 as if they were reports from any client connected to its ports, and generates a consolidated proxy report to the designated router. In this way, the router receives a single consolidated report from that entire subnet.

After the switches learn which ports are requesting access to the IP Multicast stream, all other ports not responding to the queries are blocked from receiving the IP Multicast (Figure 5).

Figure 5 BayStack 380 Switch filtering IP multicast streams (2 of 2)

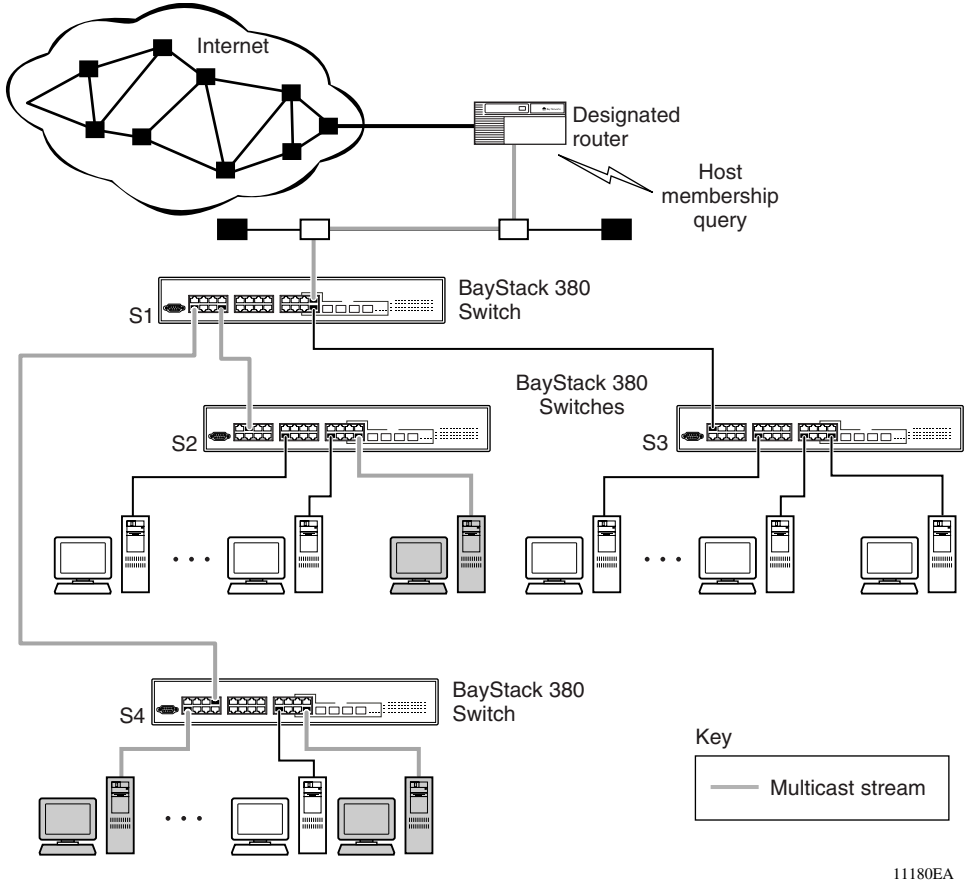

The consolidated proxy report generated by the switch remains transparent to layer 3 of the International Organization for Standardization, Open Systems Interconnection (ISO/OSI) model. (The switch IP address and MAC address are not part of proxy report generation.) The last reporting IGMP group member in each VLAN represents all of the hosts in that VLAN and IGMP group.

IGMP snooping configuration rules

The IGMP snooping feature operates according to specific configuration rules. When configuring your switch for IGMP snooping, consider the following rules that determine how the configuration reacts in any network topology:

- A port that is configured for port mirroring cannot be configured as a static router port.
- If a MultiLink Trunk member is configured as a static router port, all of the MultiLink trunk members are configured as static router ports. Also, if a static router port is removed, and it is a MultiLink Trunk member, all MultiLink trunk members are removed as static router port members, automatically.
- Static router ports must be port members of at least one VLAN.
- The IGMP snooping feature is not STP-dependent.
- The snooping field must be enabled for the proxy field to have any valid meaning.
- Static router ports are configured per VLAN and per IGMP Version.

Note: Because IGMP snooping is set up per VLAN, all IGMP changes are implemented according to the VLAN configuration for the specified ports.

Single fiber fault detection

Single Fiber Fault Detection (SFFD) allows remote fault detection on gigabit Ethernet fiber ports. When a partial fiber break occurs, data is lost on one side of a link. SFFD detects this error condition and causes the port that is losing data to go down. This stops the loss of data.

The Single Fiber Fault Detection feature is enabled on a port by port basis for the BayStack 380-24T and BayStack 380-24F.

When a port with SFFD enabled detects link failure, it signals the linked device. As soon as the linked device receives this signal, the port on that device stops transmitting data to the SFFD-enabled port. Once a link is repaired, the ports will recover automatically.

Single Fiber Fault Detection (SFFD) has the following requirements and limitations:

- SFFD must be implemented on both sides of a link
- SFFD is enabled on a per-port basis
- By default, SFFD is not enabled
- SFFD takes approximately 50 seconds to detect a fault

SFFD is only available from the command line interface (CLI). SFFD uses the following CLI commands.

show sffd

The `show sffd` command displays the current SFFD configuration for all ports where the SFFD feature is applicable. The syntax of the `show sffd` command is:

```
show sffd
```

The `show sffd` command is in the privExec command mode.

The `show sffd` command has no parameters or variables.

Figure 6 displays sample output from the `show sffd` command.

Figure 6 show sffd command output

```
BS380#show sffd
Port  SFFD Mode
----  ---------
17    Disabled
BS380#
```

sffd enable

The `sffd enable` command enables the SFFD feature on specified ports. The syntax of the `sffd enable` command is:

`sffd [port <portlist>] enable`

The `sffd enable` command is in the interface config mode.

Table 3 describes the variables and parameters for the `sffd enable` command.

Table 3 sffd enable command parameters and variables

Parameters and variables	Description
port <portlist>	Enables SFFD for specified port or ports. If left blank, the system will use the port number specified in the interface command.

Figure 7 displays sample output from the `sffd enable` command.

Figure 7 sffd enable command output

```
BS380(config-if)#sffd enable
BS380(config-if)#show sffd
Port  SFFD Mode
----  ---------
17    Enabled
BS380(config-if)#
```

no sffd enable

The `no sffd enable` command disables the SFFD feature on specified ports. The syntax for the `no sffd enable` command is:

```
no sffd [port <portlist>] enable
```

The `no sffd enable` command is in the interface config mode.

Table 4 describes the variables and parameters for the `no sffd enable` command.

Table 4 no sffd enable command parameters and variables

Parameters and variables	Description
port <portlist>	Disables SFFD for specified port or ports. If left blank, the system will use the port number specified in the interface command.

Figure 8 displays sample output from the `no sffd enable` command.

Figure 8 no sffd enable command output

```
BS380(config-if)#no sffd enable
BS380(config-if)#show sffd
Port  SFFD Mode
----  ---------
17    Disabled
BS380(config-if)#
```

default sffd enable

The `default sffd enable` command returns the SFFD configuration for the specified port or ports to the default factory setting, which is disabled. The syntax for the `default sffd enable` command is:

`default sffd [port <portlist>] enable`

The `default sffd enable` command is in the interface config mode.

Table 5 describes the variables and parameters for the `default sffd enable` command.

Table 5 default sffd enable command parameters and variables

Parameters and variables	Description
port <portlist>	Returns SFFD configuration for specified port or ports to the factory default, which is disabled. If left blank, the system will use the port number specified in the interface command.

Figure 9 displays sample output from the `default sffd enable` command.

Figure 9 default sffd enable command output

```
BS380(config-if)#default sffd enable
BS380(config-if)#show sffd
Port  SFFD Mode
----  ---------
17    Disabled
BS380(config-if)#
```

ASCII configuration generator

The ASCII Configuration Generator (ACG) allows the configuration settings of the switch to be saved to an external ASCII configuration file made up of a series of CLI commands. This editable ASCII configuration file can then be uploaded to a switch from an external file server.

Note: You must reset the switch to the factory default settings before uploading the ACG-generated ASCII configuration file. Resetting the switch will cause loss of connectivity and loss of the current configuration of the switch. Refer to *Using the BayStack 380-24T 10/100/1000 Switch* and *Using the BayStack 380-24F Gigabit Switch* for information on resetting the switch to factory default settings.

The ASCII configuration file contains configuration settings for the following network management applications:

- Core applications (system information, topology, etc.)
- Multilink Trunking
- Port configuration
- Partial Spanning Tree configuration, including configuration of port priority and path cost
- VLAN configuration

The ACG is only available from the command line interface (CLI). This section discusses the following new or enhanced CLI commands used for the ASCII Configuration Generator:

show running-config command

The `show running-config` command displays the current running configuration. The syntax for the `show running-config` command is:

```
show running-config
```

The `show running-config` command is in the privExec command mode.

Note: The `show running-config` command is not available the logged on user has read-only access.

The `show running-config` command has no parameters or variables.

Figure 10 displays sample output from the `show running-config` command.

Figure 10 show running-config command output

```
BS380#show running-config
enable
config t
mac-address-table aging-time 300
autotopology
snmp-server authentication-trap enable
snmp-server contact "SysAdmin"
snmp-server name "BS380"
snmp-server location "Lab"
snmp-server community "public" ro
snmp-server community "private" rw
--More--
```

copy running-config command

The `copy running-config` command stores the current configuration as an ASCII file on the TFTP server. The syntax for the `copy running-config` command is:

```
copy running-config tftp [address <A.B.C.D>] filename <WORD>
```

Note: The `copy config` command will copy a binary configuration file to the TFTP server. To store the configuration as an ASCII file, you must use the `copy running-config` command.

The `copy running-config` command is in the privExec command mode.

Table 6 describes the parameters and variables for the `copy running-config` command.

Table 6 copy running-config command parameters and variables

Parameters and variables	Description
address <A.B.C.D>	Specifies the TFTP server IP address; enter in dotted-decimal notation.
filename <WORD>	Specifies the name of the existing ASCII configuration file on the TFTP server. This file must be read/write enabled.

Figure 11 displays sample output from the `copy running-config` command.

Figure 11 copy running-config command output

```
BS380#copy running-config tftp address 134.177.118.56 filename config.txt
%Contacting TFTP host: 134.177.118.56.
%ACG Configuration file successfully written.
BS380#
```

configure network command

The `configure network` command loads the ASCII configuration file from an external TFTP server. The syntax for the `configure network` command is:

```
configure network [address <A.B.C.D>] [filename <WORD>]
```

The `configure network` command is in the PrivExec mode, Global configuration mode, and Interface configuration mode.

Table 7 describes the parameters and variables for the `configure network` command.

Table 7 configure network command parameters and variables

Parameters and variables	Description
address <A.B.C.D>	Specifies the TFTP server IP address; enter in dotted-decimal notation.
filename <WORD>	Enter the name of the ASCII configuration file you want to copy from the TFTP server.

Figure 12 displays sample output from the `configure network` command.

Figure 12 configure network command output

```
BS380#configure network address 134.177.118.56 filename config.txt
Downloading Config File [|]
BS380#enable
Downloaded file successfully, executing . . .
BS380#config t
Enter configuration commands, one per line.  End with CNTL/Z.
BS380(config)#mac-address-table aging-time 300
BS380(config)#autotopology
BS380(config)#snmp-server authentication-trap enable
BS380(config)#snmp-server contact "HCS lab"
BS380(config)#snmp-server community "public" ro
BS380(config)#snmp-server community "private" rw
BS380(config)#ip bootp server disable
BS380(config)#ip default-gateway 134.177.150.1
BS380(config)#ip address 134.177.150.79
BS380(config)#ip address netmask 255.255.255.0
BS380(config)#no auto-pvid
% AutoPVID already disabled.
BS380(config)#vlan mgmt 1
BS380(config)#vlan name 1 "VLAN #1"
BS380(config)#vlan members remove 1 ALL
BS380(config)#vlan members 1 ALL
BS380(config)#vlan members 2 1-12
BS380(config)#$ed-frame disable filter-untagged-frame disable priority 0
BS380(config)#$ enable proxy enable robust-value 2 query-interval 125
BS380(config)#$ enable proxy enable robust-value 2 query-interval 125
BS380(config)#vlan mgmt 1
BS380(config)#spanning-tree priority 8000
BS380(config)#spanning-tree hello-time 2
BS380(config)#spanning-tree max-age 20
BS380(config)#spanning-tree forward-time 15
BS380(config)#interface FastEthernet ALL
BS380(config-if)#spanning-tree port 1-24 learning normal
BS380(config-if)#exit
BS380(config)#no mlt
BS380(config)#mlt 1 name "Trunk #1"
BS380(config)#mlt 2 name "Trunk #2"
BS380(config)#mlt 3 name "Trunk #3"
BS380(config)#mlt 4 name "Trunk #4"
BS380(config)#mlt 5 name "Trunk #5"
BS380(config)#mlt 6 name "Trunk #6"
BS380(config)#interface FastEthernet ALL
BS380(config-if)#no shutdown port 1-24
BS380(config-if)#snmp trap link-status port 1-24 enable
BS380(config-if)#speed port 1-24 auto
BS380(config-if)#duplex port 1-24 auto
BS380(config-if)#exit
```

configure network load-on-boot command

The `configure network load-on-boot` command is used to configure the switch to automatically download a configuration file when you reboot the switch. The syntax for the `configure network load-on-boot` command is:

```
configure network load-on-boot {disable|use-bootp|
use-config} [address <A.B.C.D>] filename <WORD>
```

The configure network load-on-boot command is in the PrivExec mode, Global configuration mode, and Interface configuration mode.

Table 8 describes the parameters and variables for the `configure network load-on-boot` command.

Table 8 configure network load-on-boot command parameters and variables

Parameters and variables	Description
{disable|use-bootp|use-config}	Specifies the settings for automatically loading a configuration file when the system boots: • disable—disables the automatic loading of the configuration file • use-bootp—specifies using the BootP file as the automatically loaded configuration file • use-config—specifies using the ASCII configuration file as the automatically loaded configuration file
address <A.B.C.D>	Specifies the TFTP server IP address; enter in dotted-decimal notation.
filename <WORD>	Enter the name of the ASCII configuration file you want to copy from the TFTP server.

Figure 13 displays sample output from the `configure network load-on-boot` command.

Figure 13 configure network load-on-boot command output

```
BS380#configure network load-on-boot use-config address 134.177.118.56 filename config.txt
BS380#
```

Related publications

For more information about BayStack 380, software version 3.0 and about using the BayStack 380-24T and BayStack 380-24F switches, refer to the following publications:

- *Reference for the BayStack 380-24T and 380-24F Management Software* (part number 214393-B)

  Describes how to use Device Manager software to manage the switch.

- *Reference for the BayStack 380-24T and 380-24F Command Line Interface* (part number 215207-A)

  Describes how to use Device Manager software to manage the switch.

- *Getting Started with the BayStack 380-24T Management Software* (part number 212861-A)

  Describes how to install the Java-based device level software management application.

- *Reference for the BayStack 380-24T Management Software* (part number 212862)

  Describes how to use the Java-based device level software management application.

- *Using Web-Based Management for the BayStack 380-24T 10/100/1000 Switch* (part number 212863-A)

  Describes how to use the Web-based management tool to configure switch features.

- *Using the BayStack 380-24F Gigabit Switch* (part number 214391-A)

  Describes how to install and use the BayStack 380-24F Gigabit Switch; includes instructions to use the console interface to configure the switch.

- *Installing the BayStack 380-24F Gigabit Switch* (part number 214390-A)

  Provides installation instructions for the switch in English and five other languages.

- *Reference for the BayStack 380-24F Management Software* (part number 214393-A)

  Describes how to use Device Manager software to manage the switch.

- *Using Web-Based Management for the BayStack 380-24F Gigabit Switch* (part number 214394-A)

  Describes how to use the Web-based management interface to configure and monitor switch operations.

Hard-copy technical manuals

You can print selected technical manuals and release notes free, directly from the Internet. Go to the www.nortelnetworks.com/documentation URL. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Use Adobe* Acrobat Reader* to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to Adobe Systems at the www.adobe.com URL to download a free copy of the Adobe Acrobat Reader.

How to get help

If you purchased a service contract for your Nortel Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel Networks service program, contact one of the following Nortel Networks Technical Solutions Centers:

Technical Solutions Center	Telephone
Europe, Middle East, and Africa	(33) (4) 92-966-968
North America	(800) 4NORTEL or (800) 466-7835

Technical Solutions Center	Telephone
Asia Pacific	(61) (2) 9927-8800
China	(800) 810-5000

An Express Routing Code (ERC) is available for many Nortel Networks products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to http://www.nortelnetworks.com/help/contact/erc/index.html and follow the directions on the page.