User Guide

SmartFlow

March 2006

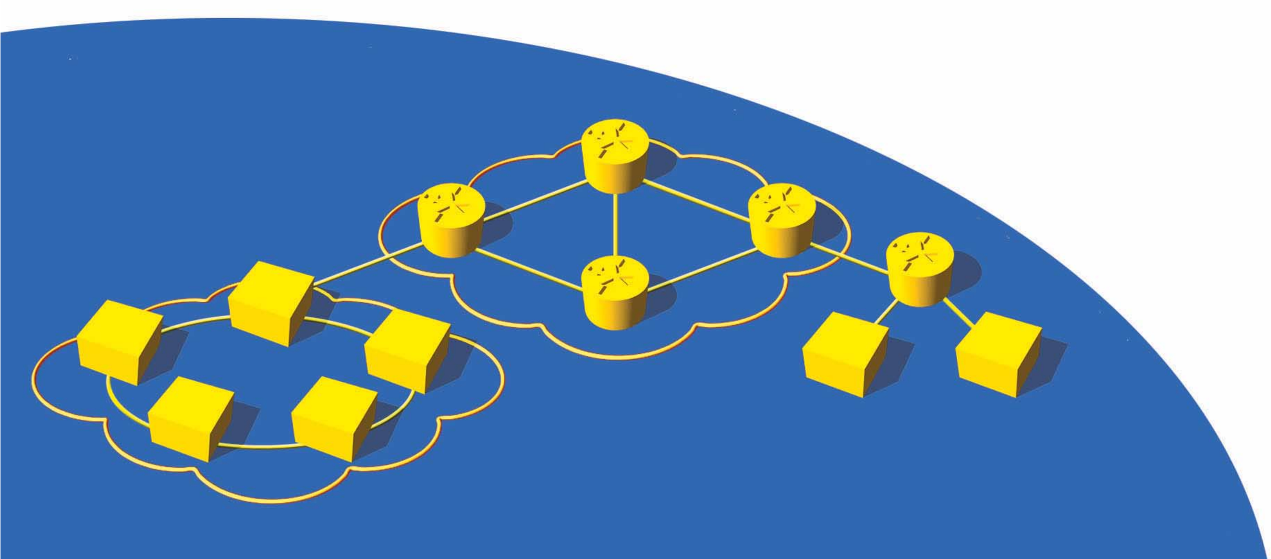

P/N 71-000622 REV A

Spirent Communications, Inc.

26750 Agoura Road
Calabasas, CA
91302 USA

Support Contacts

U.S.
E-mail: support@spirentcom.com
Web: http://support.spirentcom.com
Toll Free: +1 800-SPIRENT (+1 800-774-7368)
Phone: +1 818-676-2616
Fax: +1 818-880-9154

China
E-mail: supportchina@spirentcom.com
Web: http://support.spirentcom.com.cn
Toll Free: +1 800-810-9529 (mainland China only)
Phone: +86 10 8233 0033 (rest of the world)
Fax: +86 10 8233 0022

Europe
E-mail: support.europe@spirentcom.com
Web: http://support.spirentcom.com
Phone: +33 (0) 1 61 37 22 70
Fax: +33 (0) 1 61 37 22 51

Copyright

Limited Warranty

Spirent Communications, Inc. ("Spirent") warrants that its Products will conform to the description on the face of order, that it will convey good title thereto, and that the Product will be delivered free from any lawful security interest or other lien or encumbrance.

Spirent further warrants to Customer that hardware which it supplies and the tangible media on which it supplies software will be free from significant defects in materials and workmanship for a period of twelve (12) months, except as otherwise noted, from the date of delivery (the "Hardware Warranty Period"), under normal use and conditions.

To the extent the Product is or contains software ("Software"), Spirent also warrants that, if properly used by Customer in accordance with the Software License Agreement, the Software which it supplies will operate in material conformity with the specifications supplied by Spirent for such Software for a period of ninety (90) days from the date of delivery (the "Software Warranty Period"). The "Product Warranty Period" shall mean the Hardware Warranty Period or the Software Warranty Period, as applicable. Spirent does not warrant that the functions contained in the Software will meet a specific requirement or that the operation will be uninterrupted or error free. Spirent shall have no warranty obligations whatsoever with respect to any Software which has been modified in any manner by Customer or any third party.

Defective Products and Software under warranty shall be, at Spirent's discretion, repaired or replaced or a credit issued to Customer's account for an amount equal to the price paid for such Product provided that: (a) such Product is returned to Spirent after first obtaining a return authorization number and shipping instructions, freight prepaid, to Spirent's location in the United States; (b) Customer provides a written explanation of the defect or Software failure claimed by Customer; and (c) the claimed defect actually exists and was not caused by neglect, accident, misuse, improper installation, improper repair, fire, flood, lightning, power surges, earthquake, or alteration. Spirent will ship repaired Products to Customer, freight prepaid, within ten (10) working days after the receipt of defective Products. Except as otherwise stated, any claim on account of defective materials or for any other cause whatsoever will conclusively be deemed waived by Customer unless written notice thereof is given to Spirent within the Warranty Period. Spirent reserves the right to change the warranty and service policy set forth above at any time, after reasonable notice and without liability to Customer.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, ALL IMPLIED WARRANTIES, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, ARE HEREBY EXCLUDED, AND THE LIABILITY OF SPIRENT, IF ANY, FOR DAMAGE RELATING TO ANY ALLEGEDLY DEFECTIVE PRODUCT SHALL BE LIMITED TO THE ACTUAL PRICE PAID BY THE CUSTOMER FOR SUCH PRODUCT. THE PROVISIONS SET FORTH ABOVE STATE SPIRENT'S ENTIRE RESPONSIBILITY AND CUSTOMER'S SOLE AND EXCLUSIVE REMEDY WITH RESPECT TO ANY BREACH OF ANY WARRANTY.

Contents

About This Guide

In About this Guide...

Introduction

This manual provides information about the features and capabilities of SmartFlow and its interface. It also includes details on software installation requirements, as well as information about:

- Basic navigation
- High-level usage information such as how to set up ports and tests.
- Detailed theory of operation
- Detailed test methodologies, sample results, and how to interpret test results.

For field descriptions or detailed procedures, refer to the online Help. (The online Help also contains most of the information in this manual.)

This manual is intended for users of SmartFlow software. It is assumed that users of this manual have the following:

- Familiarity with Microsoft Windows
- Moderate familiarity with SmartBits equipment
- Intermediate level of data communications theory.

Related Manuals

The following documentation may help you understand and use SmartFlow:

- *SmartBits 600x/6000x Installation Guide*
  Describes how to set up your SmartBits chassis, and includes multi-chassis and cabling information.
- *SmartBits System Reference Manual*
  Describes the SmartBits hardware (chassis and cards/modules) including the LED display on SmartCards and modules.
- *SmartWindow User Guide*
  Describes how to perform tests using SmartWindow - the top-level software application for all SmartBits hardware.
- *Application Note #19, SmartFlow: The Right Tool for Testing QoS*
  Briefly describes how to interpret the test results in SmartFlow and how to use it to test a DUT's priority scheme, using a simple traffic pattern as an example.
- *Application Note #36, SmartBits SmartMetrics Tests: Sequence Tracking*
  Describes how subfields within the proprietary SmartBits signature field are utilized to determine frame-by-frame performance results such as latency, latency variation, etc.

- *Application Note #37, MPLS: QoS Testing with SmartFlow*
  Discusses QoS as a traffic engineering solution in Multiprotocol Label Switching (MPLS) and shows how you can test the QoS performance of your MPLS network, using TeraMetrics modules and SmartFlow.
- *SmartBits Automation for SmartFlow*
  Explains how to use the SmartBits Automation programming interface to develop test scripts for SmartFlow applications. Includes information on the Script Automation Interface (SAI), which you can use to prepare your initial test configuration. SmartBits Automation SAI enables you to write your test script in a standard text file.

Online Help

SmartFlow includes online help for the wizards, dialog boxes, tabs, menus, and common tasks. It also contains detailed procedures and field descriptions.

Online Help is available in Microsoft HTML format.

To access the online Help, select **Help** from the SmartFlow menu bar or press **F1** at any dialog box, tab, or menu.

SmartBits Hardware Handling/Cleaning Practices

SmartBits cards and modules contain electronic components that are sensitive to Electrostatic Discharge (ESD) damage. To prevent premature component failure or latent product damage, it is crucial that you handle this equipment following industry standard ESD handling practices. Refer to *Appendix B, "ESD Requirements"* for further information.

Some SmartBits equipment contains fiber optic components that are very susceptible to contamination from particles of dirt and dust. Product performance may be damaged if these components are not kept clean. Refer to *Appendix C, "Fiber Optic Cleaning Guidelines"* for proper cleaning practices for these components.

How to Contact Us

To obtain technical support for any Spirent Communications product, please contact our Support Services department using any of the following methods:

U.S.

E-mail: support@spirentcom.com
Web: http://support.spirentcom.com
Toll Free: +1 800-SPIRENT (+1 800-774-7368) (US and Canada)
Phone: +1 818-676-2616
Fax: +1 818-880-9154
Operating Hours: Monday through Friday, 06:00 to 18:00 Pacific Time

Europe

E-mail: support.europe@spirentcom.com
Web: http://support.spirentcom.com
Phone: +33 (0) 1 61 37 22 70
Fax: +33 (0) 1 61 37 22 51
Operating Hours: Monday through Thursday, 09:00 to 18:00; Friday, 09:00 to 17:00 Paris Time

China

E-mail: supportchina@spirentcom.com
Web: http://support.spirentcom.com.cn
Toll Free: +1 800-810-9529 (mainland China only)
Phone: +86 10 8233 0033 (rest of the world)
Fax: +86 10 8233 0022
Operating Hours: Monday through Friday, 09:00 to 18:00 Beijing Time

The latest versions of user manuals, application notes, and software and firmware updates are available on the Spirent Communications Customer Service Center websites at http://support.spirentcom.com and http://support.spirentcom.com.cn (China).

Information about Spirent Communications and its products and services can be found on the main company websites at http://www.spirentcom.com and http://www.spirentcom.com.cn (China).

Company Address

Spirent Communications, Inc.
26750 Agoura Road
Calabasas, CA 91302
USA

Chapter 1

SmartFlow Overview

This chapter provides an overview of SmartFlow and its features, a summary of tests, and information on frequently asked questions.

In this chapter...

What is SmartFlow?

SmartFlow is a performance analysis tool to test Layers 2, 3, and 4 on Class of Service (CoS) devices and networks built with CoS priority strategies. SmartFlow allows you to set up multiple flows of IP frames to simulate network traffic and measure the latency, frame loss, and throughput. It generates results in charts and tables that include measurements for latency, frame loss, and standard deviation of flows. You can track results by priority or by type of traffic to see the effect a prioritizing CoS device has on the network.

SmartFlow can also be used to test non-prioritizing devices by simply generating flows and then tracking and grouping them. The SmartFlow Group Wizard quickly and automatically generates any number of flows and groups of flows based on traffic pattern and priority settings that you specify in several Wizard dialog boxes.

SmartFlow is used together with a Spirent Communications SmartBits® chassis and at least two SmartMetrics or TeraMetrics® (or TeraMetrics-based) ports.

Features of SmartFlow

SmartFlow supports the following features:

- Variable, multiple source, and destination address fields *within one flow definition* to simulate a large network
- IP, TCP, UDP, ICMP, and IGMP protocols
- LAN, POS, WAN, and ATM interfaces
- Test devices with VLAN, IP TOS, and Diffserv priority schemes
- Multiple networks and VLANs per physical port
- Flows with MPLS labels, statically or dynamically defined
- Ability to test single and multiple devices
- Test results in a chart or table format that is exportable to HTML format
- Test setup exportable to a script format
- Ability to save test/card setup information and results to a file for later use
- User-configurable test duration, frame length, test load, and type of data flow
- Ability to measure throughput, frame loss, latency, and QoS reclassification for a large number of flows and groups of flows
- Generation and advertisement of a large number of BGP4 routes
- Multiple chassis connections
- Multiple users of a SmartBits chassis.

For a list of features that are new to this release, refer to the online Help.

Summary of Tests

SmartFlow contains the following tests:

- **Throughput**
  Measures the maximum rate at which frames from flows and groups sent through a device can be sent without frame loss. There is a standard Throughput test and an asymmetric Throughput test.
- **Jumbo**
  Measures frame loss, latency, and latency distribution in flows and groups of flows sent through a device and updates each of these results simultaneously. The test includes both cut-through and store-and-forward latency measurements. It also measures latency standard deviation based on latency distribution and the number of frames received that were out of sequence. It also reports bandwidth utilization.
- **Frame Loss**
  Measures the number of frames lost from flows and groups sent through a device as well as the number of frames received that were out of sequence.
- **Latency**
  Measures the minimum, maximum, and average latency of frames received on non-TeraMetrics ports plus minimum and maximum or minimum and average latency of frames received on all TeraMetrics and TeraMetrics-based ports. The test includes both cut-through and store-and-forward latency measurements. (The frames are in flows and groups of flows sent through a device.)
- **Latency Distribution**
  Measures the latency of received frames in flows and groups of flows sent through a device and sorts them into eight latency buckets. Compared to the Latency test, this test can provide a finer view of latency behavior at the Device Under Test (DUT) load tolerance limits.
- **Latency Over Time**
  Measures the overall pattern of latency (across time) for each receiving port in the test.
- **Latency SnapShot**
  Measures the latency of each received frame for a specified number of frames in flows and groups of flows sent through a device.
- **SmartTracker**
  Measures the ability of a DUT to reclassify packets based on IP precedence settings such as protocol type. It can measure other criteria relative to the VLAN tag of a frame, such as whether the VLAN tag was retained in received frames and whether it was in the correct place. You can also specify a custom field to track in a packet that is 16 bits or less.

Notes: • Latency is calculated for all received frames.

- For Latency and Latency Over Time tests, SmartFlow reports either average and maximum latency or minimum and maximum latency for all TeraMetrics and TeraMetrics-based receiving ports.

Uses of SmartFlow

SmartFlow's capacity for high port density allows you to put the DUT backbone or system under load with hundreds of ports. The main benchmarking purposes for SmartFlow are to:

- Test priority schemes.
  - Verify the priority scheme functionality of a device and how many levels it can effectively use.
  - Test *your* implementation of priority schemes.
  - Evaluate the performance at all priority levels based on different protocols.
- Determine and understand the performance limits of your devices. Then make the necessary modifications to boost its performance.
- Perform comparative testing of devices.
- Re-qualify a device after firmware upgrades.
- Determine if you are getting the bandwidth for which you paid.
- Publish test results in industry-sponsored comparative analyses and reports.
- Get your product to market more quickly.
- Advance CoS technology by illustrating its benefits with industry-standard test results.
- Test firewalls to see the number of flows the device can pass and reject.
- Test BGP4 routers.
- Test MPLS routing performance.

Who Should Use SmartFlow?

The typical users of SmartFlow are:

- Network Equipment Manufacturers (NEMs)
  (This includes developers, Quality Assurance engineers, marketing staff, and pre-sales SEs.)
- Internet Service Provider test-lab engineers
- Enterprise test-lab engineers
- Anyone who wants to deploy critical or prioritized traffic.

Why Test with SmartFlow

The increase in voice and video transmissions over IP has made prioritizing network traffic a necessity. Both multimedia and data sharing applications require efficient priority implementation in CoS network devices. SmartFlow allows you to test:

- CoS devices for DiffServ, VLAN, and TOS Precedence
- Bursty traffic
- TCP/UDP source port priority schemes
- VLAN implementations
- Basic firewall datagram forwarding
- Device access lists
- Flows between any types of ports (Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, USB, POS, WAN, or ATM)
- MPLS-enabled traffic
- IP Multicast traffic
- IPv6 traffic
- Bandwidth provisioning.

SmartFlow can help you evaluate the performance of your DUT or System Under Test (SUT) by answering such questions as:

- For priority schemes:
  - What is the impact of priority-based traffic on non-priority-based traffic?
  - Under congestion, at what point is low priority traffic dropped and how is high priority traffic treated? (Does the latency increase or decrease?)
  - How many priority queues does the device support?
  - Were flows received in sequence and were there any stray (misdirected flows)?
  - How does the DUT handle and prioritize IPv4 traffic together with IPv6 traffic?
- For voice and video data:
  - What is the line quality of the high priority traffic?
  - Does the DUT promote transmission of voice and video data?
  - Do the latency results of a port indicate that it can handle voice or video?
  - How many flows can the network handle at different line rates?

Note: Although SmartFlow does not specifically generate voice or video traffic, you can generate UDP traffic, which is one of the most common transport protocols for voice and video. SmartVoIPQoS, part of the SmartFlow Application Test Suite, does simulate voice traffic and can generate both UDP and RTP traffic.

- What is the latency and latency distribution at different priorities and loads?

The Need to Test Bursty Traffic

Network devices often produce bursts of back-to-back frames, rather than a steady load of frames. A back-to-back burst is a number of frames sent at 100% load for a specific duration. For example, in Ethernet networks, bursty traffic consists of sending frames with the minimum interframe gap.

Typically, back-to-back bursts occur at aggregate points such as an uplink between two devices. At an aggregate point, a device's buffer fills due to simultaneously receiving packets on different input ports that are destined for the same output port. The device may then send a burst of traffic from the congested output port in an attempt to rapidly empty the buffer.

How Bursty Traffic Affects Device Performance

A device's behavior can vary in the presence of bursty traffic. Latency may increase due to the additional buffering that results from receiving a burst of traffic. Throughput can also be affected by bursty traffic because a device may not be able to sustain 100% load throughout the duration of the burst.

Because normal network traffic is bursty, a device should be tested using different burst sizes. SmartFlow can create bursty test traffic based upon user-specified parameters such as the burst size, simulating real world conditions and stressing the forwarding capabilities of the device. SmartFlow has a test for throughput as well as various latency tests that help measure and examine latency.

When to Test Class of Service Devices

- When increasing multimedia unicasts and multicasts on your network.
- After network device firmware upgrades. (Do not assume that performance will improve after an upgrade.)
- After any kind of network growth.
- After you implement a priority scheme for your network or for devices that support priority.
- After you have determined the core functionality of a networking device and you now need to measure the device's scalability and performance.

Related RFCs and References

The following general references apply:

RFC 1242	Benchmarking Terminology for Network Interconnection Devices
RFC 2544	Benchmarking Terminology for Network Interconnect Devices
RFC 2285	Benchmarking Terminology for LAN Switching Devices
RFC 1700	Assigned Numbers (for IP protocol and TCP/UDP port numbers)
RFC 793	Transmission Control Protocol (TCP)
RFC 768	User Datagram Protocol (UDP)
RFC 792	Internet Control Message Protocol (ICMP)
IEEE 802.1Q	VLAN Tagging
RFC 2113	IP Router Alert Option

These references apply specifically to BGP:

RFC 1771	Border Gateway Protocol 4 (BGP-4), current version of BGP
RFC 1654	Border Gateway Protocol 4 (BGP-4), first specification now obsoleted by 1771
RFCs 1105, 1163, 1267	Describe versions of BGP prior to BGP4

These references apply specifically to MPLS:

RFC 2205	Resource Reservation Protocol (RSVP)
RFCs 2210, 2211, 2215	Describe the use of RSVP with IETF Integration Service and the Controlled Load Network element
RFCs 3031, 3032	Describes MPLS architecture and Label Stack Encoding

These references apply specifically to IP Multicast:

RFC 2432	Terminology for IP Multicast Benchmarking
IETF *Draft*	Methodology for IP Multicast Benchmarking
RFC 1112	Host extensions for IP Multicasting, specifies Version 1 of IGMP

RFC 2236	Internet Group Management Protocol, Version 2 of IGMP
RFC 2113	IP Router Alert Option

These references apply specifically to IPv6:

RFC 2373	IP Version 6 Addressing Architecture
RFC 2460	Internet Protocol, Version 6 (IPv6) Specification
RFC 2461	Neighbor Discovery for IP Version 6 (IPv6)
RFC 2462	IPv6 Stateless Address Autoconfiguration
RFC 2463	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. Neighbor discovery uses ICMPv6.
RFC 2464	Transmission of IPv6 Packets over Ethernet Networks
RFC 2473	Generic Packet Tunneling in IPv6 Specification
RFC 3056	Connection of IPv6 Domains via IPv4 Clouds

Frequently Asked Questions (FAQ)

What exactly is a "flow?"

A *flow* is a series of frames traveling from one source to one destination.

What is a SmartFlow?

A *SmartFlow* is a trackable entity that contains one or more flows. Each SmartFlow contains a series of frames from one or more sources to one or more destinations at the MAC, IP, TCP/UDP, and/or other protocol levels with their specific priority settings. (In SmartFlow documentation, the term *flow* refers to either a SmartFlow or a flow, depending on the context.)

What is a Group?

A *group* consists of one or more SmartFlows, usually with a user-defined common attribute such as priority, protocol, and traffic pattern. Groups are used to organize SmartFlow test results.

What protocols are supported in SmartFlow?

SmartFlow allows you to define flows with the IP, TCP, UDP, ICMP, or protocols over Ethernet II, POS, WAN, and ATM.

Does SmartFlow implement a TCP stack?

No, SmartFlow only allows you to define a TCP frame with a source and destination port.

What type of priorities or CoS are supported?

SmartFlow supports VLAN priority, Diffserv, IP Precedence, and MPLS labels. You can combine VLAN with IP Precedence and Diffserv. SmartFlow also allows you to enter the IP TOS (Diffserv) byte in hexadecimal form or select from a pull-down list of values.

Can SmartFlow detect and track when the TOS byte has changed?

Yes. You can use the *SmartTracker* test to track the TOS byte, as well as other prioritization bits, in test frames. Refer to *Chapter 16, "SmartTracker Test."*

How many flows per port can be defined and used?

SmartFlow supports numerous source and destination pairs from or to a single port. (The number of flows can be less depending on the number of multiples defined in each flow.) Depending on the individual card/module, the number of possible cyclic flows per port is significantly greater. See the table in *"Maximum Number of Flows per Port" on page 45*.

How do I create a large number of flows?

You can create and edit each flow one at a time or use the Group Wizard. The wizard allows you to create a large number of flows between ports in a mesh or backbone configuration or in pairs spanning all the priorities. You can also use an existing flow(s) as a "template" to set up more flows using Copy Flow function the *Smart-Flows* tab.

Can I send background traffic?

SmartFlow doesn't include a separate feature that allows you to specify background traffic to be sent along with the flows. However, you can create any flows for background purposes. You can also use other SmartBits applications to send background traffic from non-SmartFlow ports.

How many groups can a flow belong to?

A flow can belong to as many groups as you define. You can create any number of groups and have any or all flows belong to the group. The groups are intended to allow you to get results based on different priorities, source or destination addresses, or any other logical grouping.

Can groups span multiple ports?

Yes, groups are made up of user-defined flows that can be between many different ports. This is a key feature in creating groups.

How many networks per port can be used?

Up to 999 networks or less, depending on the subnet mask and the overall memory limitations of the test being run. The subnet is defined by the device's gateway address and subnet mask. You can enter multiple subnets for each port.

How are the results displayed?

The results, which are stored in a Microsoft Access database format, are displayed by flow, group, and total. The results are presented for each test graphically, in table format, and with raw data values. The graphs and tables can be cut and pasted into other applications. The graphs can be saved as .jpegs and the tables can be saved to comma delimited files (.csv), which can be read by Excel, or to .html files.

What kind of device can be tested with SmartFlow?

SmartFlow can test any forwarding device. However, it is intended to test priority-based devices. It can teach Layer 2, Layer 3, and VLAN devices and learn the MACs of Layer 3 devices. It supports ARP for each flow and can forward frames through a VLAN device.

Can I test multiple devices simultaneously?

Yes, as is the case with all Spirent Communications applications. No assumptions are made regarding the internal nature of the DUT. Thus, the DUT can be comprised of multiple devices making it an SUT.

How many ports does SmartFlow support?

The number of ports per SmartBits system determines the maximum number of ports supported. SmartFlow supports up to 16 SmartBits systems (as 16 SmartBits 6000x's, or 16 stacks consisting of either four SmartBits 2000s or one SmartBits 2000 and three SmartBits 10s). However, Latency tests support 8 SmartBits systems (or 16 with a multi-chassis extension unit) when using a CAT5 cable for synchronized chassis using the expansion in/out ports. For more information about chaining multiple chassis, see the *Installation Guide* shipped with your chassis.

Can SmartFlow be run on a SmartBits 6000x and a SmartBits 600x?

Yes. SmartFlow supports these chassis and various POS and LAN modules that run on these chassis. (The SmartBits 600x/6000x does not support WAN modules.) For specific modules supported, see *"SmartBits Hardware Requirements" on page 59*.

Can SmartFlow be used with multiple master chassis?

Yes, SmartFlow supports multiple chassis with synchronized time stamps and starts. You can connect chassis in various combinations: SmartBits 600x/6000x's with SmartBits 200/2000s or SmartBits 2000s with SmartBits 10 expansion chassis. You can also use a multi-chassis extension unit to connect multiple master (controlling) chassis. For more information, see the *Installation Guide* for your chassis type.

Does SmartFlow support the use of CDMA or GPS as external time references?

The SmartBits 600x/6000x chassis now supports the EndRun Technologies Praecis CDMA and Zyfer NanoSync GPS time reference as well as the original Symmetricom/HP 58503A/B unit. SmartFlow works with any combination of these units. No special settings need be set in SmartFlow; operation is automatic. For the SmartBits 200/2000, only the HP 58503A/B are supported.

How can I set up Multicast receiver groups in SmartFlow?

To set up receiver groups, use the IP Multicast Wizard. Select the **Multicast** tab, then right-click on the first column **MC Group/Port-VLAN** to start the wizard. Use the Wizard to associate IGMP-enabled ports (but enable them in the *Cards* tab first) with the multicast groups that you want the ports to join. Use the Wizard to configure multiple VLANs per multicast group as well.

How can I run a single-port test with SmartFlow?

In SmartFlow version 4.00 or later, you can do this easily. In the Group Wizard, when selecting port pairs, select the **One port configuration** checkbox. In versions earlier than 4.00, use this procedure: Set two ports to Read State “Active” in the *Cards* tab (though only one port is used for the test). Using the Group Wizard on the *Groups* tab, create a flow between the two active ports. Then from the *SmartFlows* tab, change the destination port of the flow that you created to match the source port. This effectively allows you to run a test on a looped-back circuit using a single port.

Which hardware supports rates per flow and how does it work exactly?

Rates per flow is supported on LAN-33xxA, XFP-373xA, and XLW-372xA modules. This feature on the *Groups* tab enables you to control rates by individual flows. It includes the Flow Rate Calculator, which you can use to convert rates from Kbits/Sec, Mbits/Sec, and Frames/Sec to %Utilization. When assigning load percentages to flows, remember that these ratios are distributed across the load percentage of the port on a per-iteration basis. Example: If flow A is set to 50%, flow B set to 25%, and flow C set to 25%, and the port load utilization percentage is set to 10% for the first iteration of a test, then flow A transmits at 5% of line speed, and flows B and C transmit at 2.5% of line speed.

When I setup TCP traffic in SmartFlow, does it actually open TCP sessions?

No, it builds TCP frames and will ARP for these flows, but it does not open actual TCP sessions. Spirent Communications offers applications that can do this. Consult your local Sales Representative for more information.

Why do I see frame loss when I use a larger burst number even though the load and length of the tests are the same?

When the burst value is set to 1, the Inter Frame Gap (IFG) between each of the packets is evenly distributed. If the number is set to something other than 1, that burst of frames is sent with the minimum legal gap between those frames and the remaining bursts of traffic are sent with Inter Burst gaps (IBG) to correct for the load.

For example, say you are transmitting 128-byte frames (including CRC) at a rate of 100 packets per second (pps). The port speed is 100Mbps and the burst value is 1. The IFG between each packet would be approximately 10ms (minus the time length of the packet with the preamble, 10,880ns) or 9.989120ms.

If the burst value is 10, there would be 10 bursts of 10 packet per burst per second. The IFG would be 96 bit times (960ns) and the IBG would be 99.882560ms. Because the packets are closer together in the burst, the burst fills the buffer faster, and it requires routing decisions to be made faster in the DUT. This can cause frame loss if the DUT can’t keep up. (This is for illustrative purposes only. Most DUTs keep up with traffic close to or at wire speeds.)

New in this Release

SmartFlow Version 5.50 contains these new features, enhancements, and changes:

- *Asymmetric Throughput*
  SmartFlow now reports throughput separately and independently for two sets of flows with a new test called Asymmetric Port Throughput. One flow group is assigned to upstream ports and another to downstream groups. The test uses a separate search algorithm (binary, step, or combo) for each set of flows. Upstream and downstream test traffic can be run simultaneously or sequentially. Upstream and downstream flow groups can contain sub-groups, i.e., one group of TCP flows and one group of UDP flows. The test configuration and results use meaningful load units: Mbps, Kbps, fps, and bps as well as % utilization.
- *Frame Size Summary Reporting*
  New frame size summary reports are available in the summary results of the Jumbo, Frame Loss, and Latency tests. These results make packet sweep testing easier, especially across multiple ports and traffic loads.
- *ATM in Flow Rate Calculator*
  ATM-transmit rate support and conversion between Ethernet and ATM have been added to the Flow Rate Calculator that is accessed from the *Tools* menu. With this tool, the number of ATM cells and frames can easily be computed based on traffic loads and frame contents.
- *IPv6 Management Control*
  The chassis can now be controlled from an IPv6-enabled Ethernet interface on a PC on which SmartFlow is running.
- *Random Frame Size Limits*
  Random frame sizes assigned to multiple flows on the *SmartFlows* tab can now be assigned a range with an upper and lower frame size by highlighting the flows and right-clicking to access this option.
- *Increase in Maximum Custom Loads*
  The maximum number of test loads per sequence index on the *Custom Test Loads Per Port (% Utilization)* dialog box (accessed from the *Test Setup>Test Iterations* tab) has been increased from 2,000 to 5,000.
- *SAI Export of Multicast*
  SmartFlow with Multicast functionality can now be exported to SAI via the *Export to SAI* option accessed from the *File* menu.
- *SAI Export*
  The new features in SmartFlow Version 5.50 can be exported to SAI via the *Export to SAI* option accessed from the *File* menu.

SPIRENT
Communications

Chapter 2

Basic Theory

This chapter provides some basic definitions such as flow and stream, and explains how cyclic and non-cyclic SmartFlows differ. It also provides a general test methodology.

In this chapter...

What is a Flow?

A *flow* in general is a series of frames traveling from one source to one destination.

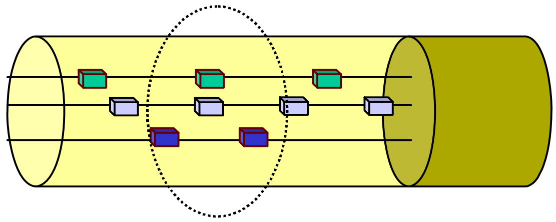

Each line represents a flow.

In SmartFlow documentation, the term *flow* can refer to either a SmartFlow or a flow, depending on the context.

What is a SmartFlow?

A *SmartFlow* is a trackable entity that contains one or more flows and provides the ability to vary addresses and/or port numbers. Each variation comprises one flow. In the application, you manipulate flows by changing the SmartFlow to which the flow belongs.

Each SmartFlow contains a series of frames from one or more sources to one or more destinations at the MAC, IP, TCP/UDP, and/or other protocol levels with their specific priority settings. There are two types of SmartFlows: cyclic and non-cyclic.

What is a Stream?

A *stream* is a template (blueprint) or pattern, based on a single protocol, for creating frames/flows. Frames in a stream are related by the basic traffic blueprint that they share, such as type (IP, UDP), frame size, and content. Streams relate directly to hardware resources on a card or module. The maximum number of streams allowed *per port* on a card or module ranges from 512 - 8,192 depending on the card type. Each non-cyclic SmartFlow uses one stream by default. For more information, see *"How a Stream Generates a Flow" on page 32*.

What is a Non-cyclic SmartFlow?

A non-cyclic SmartFlow is simply a SmartFlow that uses one or more hardware streams. It allows more combinations of addresses and/or port number (within the protocol header) variations than a cyclic SmartFlow, but each variation uses a hardware stream resource. For a comparison, see *"Maximum Number of Flows per Port" on page 45*.

What is a Cyclic SmartFlow?

A cyclic SmartFlow is a SmartFlow consisting of *one* hardware stream with incrementing fields. You can generate more address variations using a cyclic SmartFlow than a non-cyclic one. For more information, see *Table 2-6, "Maximum Number of Flows per Port," on page 45*.

What is a Group?

A *group* consists of one or more SmartFlows, usually with a user-defined common attribute such as priority, protocol, and traffic pattern. Groups are used to organize SmartFlow test results. Flows are grouped according to what you want to track in test results. Grouping flows does not affect traffic generation.

How a Stream Generates a Flow

A *stream* in SmartFlow is a frame definition (packet structure) based on one protocol. It is possible to have multiple streams that are identical except for the signature field. A stream emulates a virtual host. The stream definition is used to generate multiple test frames of the same protocol type. *Figure 2-1* shows how various types of SmartFlows use streams.

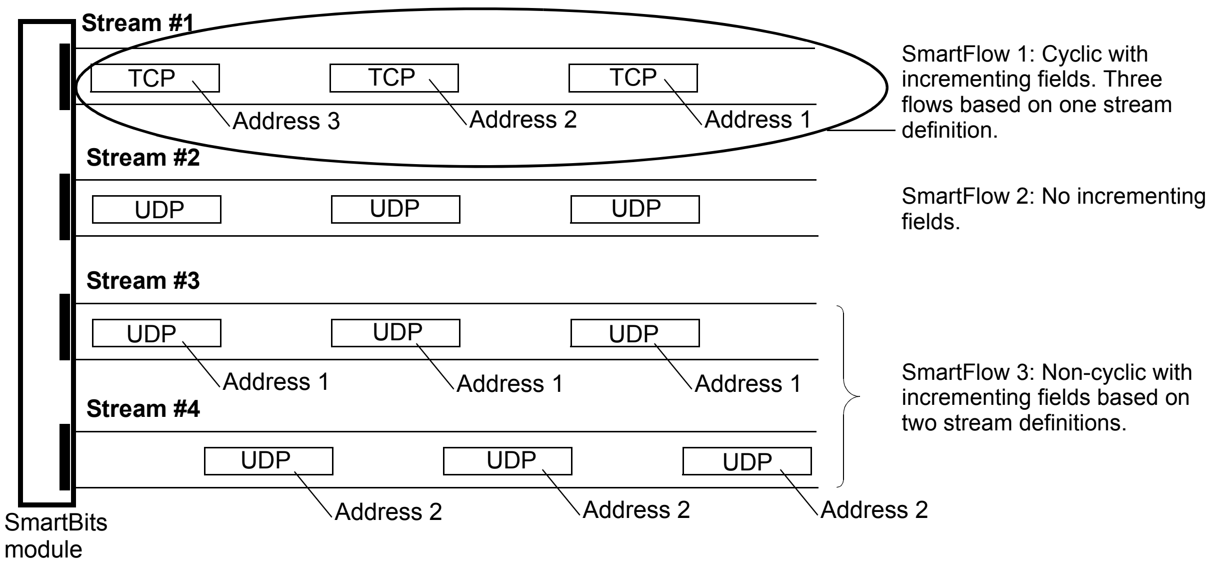

Figure 2-1. How Cyclic and Non-cyclic SmartFlows Utilize Streams

When you make an incrementing SmartFlow cyclic, one stream definition is used to generate multiple flows with varying field values (in *Figure 2-1*, **SmartFlow 1**). Here *flows* denotes frames traveling from one source to one destination, not SmartFlows.

A SmartFlow that contains no varying fields consists of a single stream definition. (**SmartFlow 2**).

A non-cyclic SmartFlow that has incrementing fields consists of multiple streams (**SmartFlow 3**). The number of stream definition is determined by the variable count.

For more information about cyclic flows, see *"Understanding Cyclic vs. Non-cyclic SmartFlows" on page 34*.

Figure 2-2 on page 33 shows one stream in a cyclic SmartFlow.

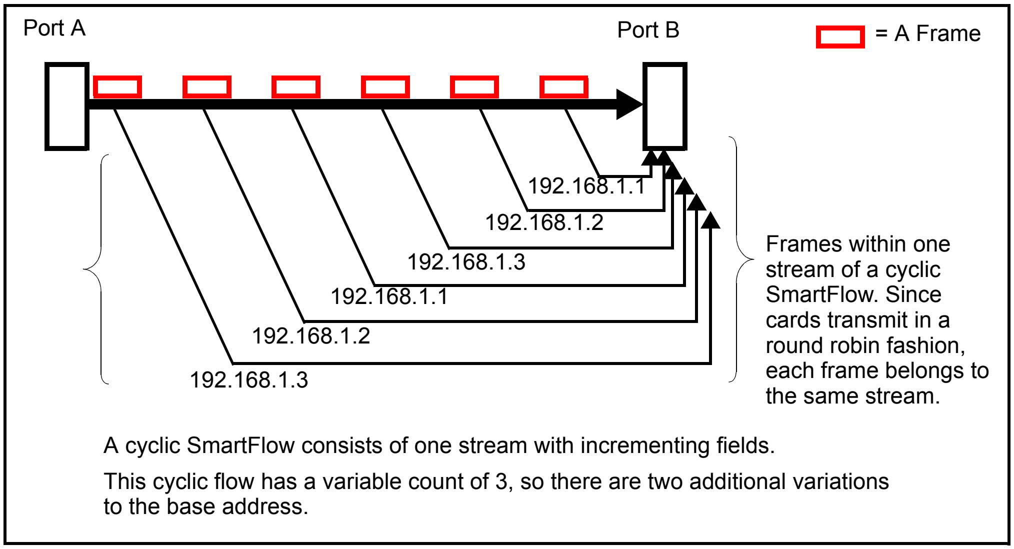

Figure 2-2. One Stream in a Cyclic Flow

Understanding Cyclic vs. Non-cyclic SmartFlows

A cyclic SmartFlow is a SmartFlow consisting of *one* stream definition with one or more incrementing fields. In cyclic flows, SmartFlow increments the fields in the frame a specified number of times. Depending on card type, the incrementing (varying) fields may be source or destination IP or MAC addresses, TCP/UDP port numbers, or VLAN ID and Priority. This way of generating varying test frames is termed *cyclic* because, after SmartFlow has varied a field the specified number of times, it reverts to the base (starting) value of each incrementing field and begins again.

When to Use Cyclic SmartFlows

Cyclic flows permit one or two varying fields, and they allow a very large number of flows (up to 65,536 per SmartFlow). Non-cyclic flows permit any number and combination of variable fields, but a much smaller number of flows. In general, use cyclic flows when you want a large number of flows. Use non-cyclic flows when you want more combinations of address and/or port number variations and a smaller number of flows is acceptable. In addition, use non-cyclic flows when you want statistics on individual flows rather than on flows by group. With cyclic SmartFlows, SmartFlow provides group totals based on each hardware stream, but it does not show results for each flow.

See the next section *"Comparison of Cyclic and Non-Cyclic SmartFlows"* for a more detailed comparison.

Notes:

- The waterfall pattern of traffic in the *Group Wizard* does not support cyclic flows. However, you can manually change the flows to cyclic on the *Traffic* tab.
- SmartFlow currently does not support cyclic flows with IPv6 traffic.

Comparison of Cyclic and Non-Cyclic SmartFlows

Either method of generating a SmartFlow has both benefits and restrictions, as shown in *Table 2-1*.

Table 2-1. Comparison of Cyclic and Non-cyclic SmartFlows

Feature	Single (cyclic) SmartFlow	Multiple (Non-cyclic) SmartFlows
Number of Flows Generated	Generates the number of address, port, or VLAN variations specified in the *Number of Flows* field within one SmartFlow.	Generates the number of SmartFlows (each with a different address) specified in the *Number of Flows* field.
Flow Limit Allowed	For information on cyclic flow limits, refer to *Table 2-6, "Maximum Number of Flows per Port," on page 45*.	The quantity is limited to the card limit on flows per port, as stated in *Table 2-6, "Maximum Number of Flows per Port," on page 45*.
CPU Usage	This allows you to generate a large number of streams with only one SmartFlow, so less resources are used.	Uses more resources since you are actually generating more individual flows to be sent and tracked.
Test Results	Only the single base SmartFlow stream is tracked and appears in results, not the individual cyclic flows.	Each SmartFlow appears in test results. You can more easily track stray frames since there are multiple flows. However, if you vary a field within the flow, the variations do not appear in results.
Variation Count (how many unique patterns)	Range of 2 - 65,535.	Range of 2 to the maximum number of streams that the card can support (e.g., 2047 for LAN-3325).
Variation Step (amount of increase between each pattern)	Limited to 1.	Limited to 1.
Variable Fields	IP and MAC source and destination addresses, TCP/UDP port numbers, and VLAN ID and Priority. See *"Card/Module Support for Cyclic Flows" on page 42* for more information. (Limited to two variable fields.)	IP and MAC source and destination addresses, TCP/UDP port numbers, and VLAN ID and Priority. (There is no limit to the number of variable fields.)

Cyclic vs. Non-Cyclic Incrementing of Flows

Incrementing fields in a cyclic SmartFlow work similarly to those in non-cyclic SmartFlows: The variable fields in the cyclic flow or SmartFlow are incremented the specified number of times. But with cyclic flows, the order of incrementing is different from that in a non-cyclic SmartFlow.

Figure 2-3 is an example, showing how IP addresses are incremented and transmitted in both cyclic and non-cyclic flows.

Let us say a card/module is configured to transmit three flows: two non-cyclic SmartFlows and one cyclic SmartFlow. The cyclic SmartFlow has the *Number of Flows* field set to 4. As the card goes through its table of flows to transmit, the address in the cyclic SmartFlow is incremented for three transmission cycles, then returns to the base address and begins incrementing again.

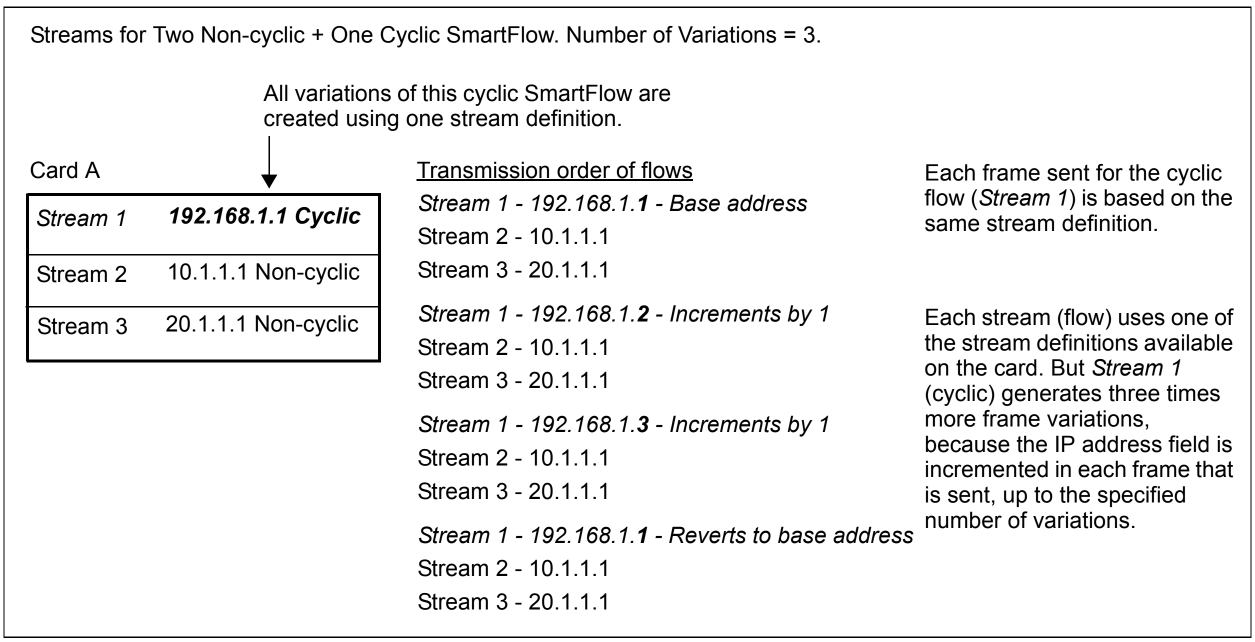

Figure 2-3. How Cyclic and Non-cyclic Flows are Transmitted

Both cyclic and non-cyclic flows use separate streams on a port. When you generate multiple non-cyclic SmartFlows, each flow uses one stream. In contrast, when you increment variable fields in a cyclic SmartFlow, each variation is performed on the same stream.

Multiple Non-cyclic SmartFlows with No Variation

For example, let us say four SmartFlows have these IP addresses with no variation:

Multiple non-cyclic SmartFlows with Number of Flows (Variable Count) = **1** (no variation)

Flow 1: 10.10.1.1

Flow 2: 10.10.2.1

Flow 3: 10.10.3.1

Flow 4: 10.10.4.1

Single Non-cyclic SmartFlow with Variation

Here is how addresses are incremented with one non-cyclic SmartFlow, whose base address is 10.10.1.1:

One Non-cyclic Flow with Number of Flows (Variable Count) = 4

10.10.1.1
10.10.1.2
10.10.1.3
10.10.1.4
10.10.1.1

This single, cyclic SmartFlow has 4 streams. Each time the port transmits this flow the IP address is incremented by one.

The fifth transmission returns to the starting address again.

Cyclic SmartFlows with Variation

For cyclic SmartFlows, SmartFlow increments the x.y.z bytes in an 1.x.y.z address, starting from the .z (least significant byte).

Here is how addresses are incremented with two cyclic SmartFlows, whose base addresses are 10.10.1.1 and 20.10.1.1:

Two Cyclic SmartFlows with Number of Flows (Variable Count) = 4

Flow 1: 10.10.1.1

Flow 2: 20.10.1.1

This is the end of the first transmission.

Flow 1: 10.10.1.2

Flow 2: 20.10.1.2

This is the end of the next transmission.

Flow 1: 10.10.1.3

Flow 2: 20.10.1.3

This is the end of the next transmission.

Flow 1: 10.10.1.4

Flow 2: 20.10.1.4

Cyclic and Non-cyclic SmartFlows with Variation

With non-cyclic flows, there is one stream definition for each pattern of incrementing fields. In addition, the *order* of stream transmission as fields increment is different from that of cyclic flows. *Figure 2-4* is an example. It shows how addresses are incremented and frames are transmitted when there are two cyclic SmartFlows and non-cyclic SmartFlow. Notice that the non-cyclic Flow 1 accounts for two-thirds of all traffic, whereas Flows 2 and 3 (cyclic) account for only one-sixth.

One Non-Cyclic, Two Cyclic SmartFlows. Number of Flows (Variable Count) = 4

Flow 1: 192.168.1.1
192.168.1.2
192.168.1.3
192.168.1.4

Flow 1 is a non-cyclic. It uses 4 hardware streams. The first 4 frames from the port are for Flow 1.

Flow 2: 10.10.1.1

Flows 2 and 3 are cyclic.

Flow 3: 20.10.1.1

This is the end of the first transmission.

Flow 1: 192.168.1.1
192.168.1.2
192.168.1.3
192.168.1.4

Flow 2: 10.10.1.2

The addresses of cyclic SmartFlows increment once each time the port transmits that flow.

Flow 3: 20.10.1.2

This is the end of the next transmission.

Flow 1: 192.168.1.1
192.168.1.2
192.168.1.3
192.168.1.4

Flow 2: 10.10.1.3

Flow 3: 20.10.1.3

This is the end of the next transmission.

Flow 1: 192.168.1.1
192.168.1.2
192.168.1.3
192.168.1.4

Flow 2: 10.10.1.4

Flow 3: 20.10.1.4

Figure 2-4. Example of Incrementing Fields in Cyclic and Non-cyclic Flows

How SmartFlow Schedules Flow Transmission

The order in which non-cyclic flows are transmitted depends on whether the flows are using port-based or flow-based rates. For more information about the different ways to vary loads in a test, see *"Varying Test Loads" on page 253*.

If no rate has been set on a per flow basis for any of the flows associated with a port (**N/A** is displayed on the *Groups* tab), flows are transmitted from that port in a round-robin fashion. If each flow uses the same frame size, the rate for each flow is evenly divided among the flows based on the overall port rate for a given test iteration. If each flow uses a different frame size, the rate for each flow is proportional to its frame size.

For example, let us say port 1 has four flows all with the same frame length: F1, F2, F3 and F4. Each flow has no rate assigned to it (**N/A** is displayed on the *Groups* tab in the *% / Custom* field.) *Table 2-2* shows the order in which the port transmits the flows during one test iteration.

Table 2-2. Transmission Order of Flows with Rate Based on Port Rate

Sequence of Transmission	Packet from Flow #
1	F1
2	F2
3	F3
4	F4
5	F1
6	F2
7	F3
8	F4

Flow Scheduling on TeraMetrics and TeraMetrics-based Cards

TeraMetrics and TeraMetrics-based module firmware attempts to maintain the round-robin transmission order of non-cyclic flows, as shown in *Table 2-2*. However, if you specify rates on a *per flow* basis, instead of round-robin, the firmware uses a ratio of each flow's rate to another to schedule flow transmission.

When rates per flow are used, the rates are not deterministic and transmission scheduling varies based on the following:

- Frame size for each flow on the port
- The ratio of the rates for each flow on the port
- The size of the card rate-based scheduler table (up to 8,192 entries).

Note: The standard TeraMetrics scheduler may adjust requested Tx rates on a per-flow basis. This may be noticeable at higher link speeds. (For more information about the scheduler, refer to the *SmartBits System Reference* manual.) Use the advanced TeraMetrics scheduler to maximize deterministic rates and maintain the desired scheduling.

For example, assume that port 1 has four non-cyclic flows: F1, F2, F3 and F4. Let us say also that each flow has the same frame length. *Table 2-3* shows the rate assigned to each flow when the test load is 10%.

Table 2-3. Example of Rate Ratios for Flows with Rates per Flow

Flow #[1]	% / Custom	Flow Rate Ratio	Actual Rate for Flow
F1	10%	1	1%
F2	20%	2	2%
F3	40%	4	4%
F4	30%	3	3%

1 Represents the order in which flows appear on the *SmartFlows* or *Groups* tab. It is not necessarily the order in which the flows are transmitted.

Based on the example in *Table 2-3 on page 40*, *Table 2-4* shows the order in which the port actually transmits the flows during one test iteration.

Table 2-4. Transmission Order of Non-cyclic Flows with Rates per Flow

Sequence of Transmission	Packet from Flow #
1	F1
2	F2
3	F3
4	F4
5	F2
6	F3
7	F4
8	F3
9	F4
10	F3

Cyclic Flows and Rate Scheduling

The scheduling principles that apply to non-cyclic flows also apply to cyclic flows. But when you assign a rate to a cyclic flow, that rate is evenly distributed among all of the cyclic variations. Each time the port transmits a cyclic flow, the next variation is applied to the transmission (unless the *Stutter count* field on the *SmartFlows>Traffic* tab is enabled). You cannot assign a rate per flow for each variation of a cyclic flow, but the rate that you assign applies to the base flow. For more information, see *"Cyclic vs. Non-Cyclic Incrementing of Flows" on page 36* and *"Rates per Flow" on page 264*.

Card/Module Support for Cyclic Flows

SmartBits cards and modules support incrementing fields in cyclic flows in different combinations. *Table 2-5* is a summary of these combinations.

Notes:

- With cyclic flows, the maximum number of variable fields is two.
- Not all frame sizes are supported for all the combinations in *Table 2-5*.
- For cyclic flows generated from the ML-7710 and LAN-3101A/B, the maximum achievable line rate may vary depending on frame size and the cyclic flow configuration. To assure maximum line rate, use a frame length of at least 144 bytes (without CRC).
- If you enable *Cyclic ARP requests* on the *Learning* tab, the following combinations are allowed: Src IP, Dst IP, Src/Dst IP, and Src MAC/Src IP. The following fields are not allowed: Dst MAC, Dst TCP/UDP Port.
- You cannot enable *Cyclic ARP requests* on the *Learning* tab and cyclic *VLAN ID/Priority*. Cyclic ARP does not support cyclic VLANs.

Table 2-5. Card/Module Support for Cyclic Flows

	ML-771x[1]	LAN-31xxA	POS-3500/ POS-3502A	LAN-3201x	POS-3505A POS-3504A POS-351xA	XLW-372xA XFP-373xA LAN-33xxA ATM-3451A ATM-3453A	WN-3445A
Src MAC						X[2]	
Dst MAC						X[2]	
Src/Dst MAC		X				X[2]	
Src IP	X	X	X	X	X	X	X
Dst IP	X	X	X	X	X	X	X
Src Port					X	X	
Dst Port					X	X	
Src/Dst IP		X		X	X	X	
Src/Dst Port					X	X	
Src MAC/ Src IP	X	X	X	X		X	
Src MAC/ Src Port						X	

Table 2-5. Card/Module Support for Cyclic Flows (continued)

	ML-771x[1]	LAN-31xxA	POS-3500/ POS-3502A	LAN-3201x	POS-3505A POS-3504A POS-351xA	XLW-372xA XFP-373xA LAN-33xxA ATM-3451A ATM-3453A	WN-3445A
Dst MAC/ Dst IP	X	X	X	X		X	
Src MAC/ Dst IP			X	X		X	
Dst MAC/ Src IP			X	X		X	
Dst MAC/ Src Port						X	
Src Port/ Dst IP					X	X	
Src Port/ Src IP					X	X	
Dst Port/ Dst MAC						X	
Dst Port/ Src MAC						X	
Dst Port/ Dst IP					X	X	
Dst Port/ Src IP					X	X	
VLAN ID						X[3]	
VLAN Priority						X	
VLAN ID/ VLAN Pri						X	
VLAN ID/ Src MAC						X	
VLAN ID/ Dst MAC						X	
VLAN ID/ Src IP						X	

Table 2-5. Card/Module Support for Cyclic Flows (continued)

	ML-771x[1]	LAN-31xxA	POS-3500/ POS-3502A	LAN-3201x	POS-3505A POS-3504A POS-351xA	XLW-372xA XFP-373xA LAN-33xxA ATM-3451A ATM-3453A	WN-3445A
VLAN ID/ Dst IP						X	
VLAN ID/ Src Port						X	
VLAN ID/ Dst Port						X	
VLAN Pri/ Src MAC						X	
VLAN Pri/ Dst MAC						X	
VLAN Pri/ Src IP						X	
VLAN Pri/ Dst IP						X	
VLAN Pri/ Src Port						X	
VLAN Pri/ Dst Port						X	

1 For these cards, frame sizes less than 144 bytes (without CRC) will not support wire rate. The ML-7710/11 supports only the following combinations: *Src IP/Src MAC* and *Dst IP/Dst MAC*.
2 ATM modules support only these fields, and only when the encapsulation mode is set to **LLC Bridged IP**.
3 VLAN fields are supported only by TeraMetrics-based LAN modules.

Maximum Number of Flows per Port

Table 2-6 shows the maximum number of hardware streams (non-cyclic SmartFlows) and variations (cyclic SmartFlows) allowed per transmitting port, for each SmartBits card and module.

Note: *LAN-31xx* module: If you have enabled jumbo or oversized frames (up to 16,388 bytes), the maximum number of streams is reduced to 128.

Table 2-6. Maximum Number of Flows per Port

Card/Module	Maximum Number Streams (Non-cyclic Flows) per Port	Maximum Number Cyclic Flows per Port[1]
ATM-3451Ax ATM-3453Ax	8,191	536,797,185
LAN-31xxA LAN-3101B	999	65,469,465
LAN-3201B/C	7,000	524,280,000
LAN-330xA LAN 3310/LAN-3311A[2]	511	33,488,385
LAN-3306 LAN-332xA	2047	134,150,145
XLW-3720A/XLW-3721A	511	33,488,385
XFP-3730A/XFP-3731A	2047	134,150,145
ML-7710	999	65,469,465
ML-5710	999	N/A
POS-3500A/POS-3500B POS-3502A	8,000	524,280,000
POS-3504As/POS-3505As[2] POS-3510A/POS-3511A[2] POS-3518A/POS-3519A	511	33,488,385

1 The maximum number of cyclic flows allowed is calculated by: **maximum number of streams *X* 65,535**—where 65,535 is the maximum number of flows allowed per stream. These maximums are theoretical. The actual limit depends on system resources and the 1-gigabit size limit on the database.
2 If the Non-XD TeraMetrics-based modules transmit IPv6 TCP frames, the maximum number of flows is 255.

About SmartCards and Modules

SmartCards and modules are custom-designed Printed Circuit Boards (PCBs) that fit within a SmartBits chassis to generate, capture, and track network packet data.

SmartCards consist of one PCB and fit into the SmartBits 10 and SmartBits 200/2000 chassis.

Modules consist of one PCB and fit into the SmartBits 600x and SmartBits 6000x chassis. They have a higher port density than do SmartCards. In the SmartBits 6000x, two modules fit into each slot.

The term *card* is used to refer to any SmartCard or module for SmartBits systems.

Notes:

- The term *SmartCard* applies only to the SmartBits 200/2000 chassis. The term *module* applies only to the SmartBits 600x/6000x chassis.
- *SmartModules* and the SmartBits *6000* chassis in which they are used are no longer manufactured; however, SmartFlow still supports them.

See *"SmartBits Hardware Requirements" on page 59* for a list of the TeraMetrics and TeraMetrics-based modules and SmartMetrics that are supported in SmartFlow.

Frame Sizes and Padding

In SmartFlow, displayed frame sizes include CRC, but they do not include any padding in the size.

SmartFlow calculates the bit rate for flows using the preamble, Layer 2 header, payload, CRC, and IFG (96 bits).

Undersize Ethernet TeraMetrics-based Frames

For any Ethernet Layer 2 frame smaller than 64 bytes, TeraMetrics and TeraMetrics-based transmitting ports add padding to make the frame 60 bytes (64 with CRC). Ports add padding to undersized frames to:

- Ensure that the DUT does not drop the frames in cases where it does not add padding.
- Successfully run a back-to-back test.

For example, if you enter a frame size of 56 (the minimum Ethernet frame size), the port transmits a 64-byte frame (with CRC). However, at Layer 3, the *IP Total Length (tot)* field in the frame remains at 38.

TeraMetrics and TeraMetrics-based receiving ports are capable of tracking padded frames for results.

Undersize Ethernet Non-TeraMetrics Frames

Non-TeraMetrics-based Ethernet ports do not pad transmitted frames that are less than 64 bytes (with CRC). If the receiving port is not TeraMetrics-based, these undersize frames show as lost frames in results, even if the DUT pads the frames, because SmartFlow is not able to track the signature field.

For a complete list of minimum and maximum frame sizes allowed in SmartFlow, see *“Allowable Frame Size Ranges by Card Type” on page 500*.

General Test Methodology

SmartBits ports used in a test behave as nodes on a network segment. SmartFlow allows you to set up multiple unicast streams within each flow as well as multiple flows and networks for each port. You can group together flows that you would like to view as an aggregate in test results.

SmartFlow embeds a signature tag (field) into all test frames to distinguish it from real network traffic. Each signature field contains the following information:

- Sending chassis and port number
- Flow number
- Sequence number of the packet in the flow
- Timestamp of when the frame was transmitted and received.

Notes:

- SmartFlow recognizes only packets whose signature field is intact (still resides before the checksum) upon receipt. For example, assume a port sends a 64-byte packet (including VLAN tag and CRC) to a non-VLAN tagged port, and the DUT "pads" the packet before the CRC. SmartFlow counts the packet as lost because the signature field is not immediately before the checksum.
- If you run the SmartTracker test, the alternate key is used instead of the signature field.

Latency Results

SmartFlow provides latency results in two forms. For testing *bit-forwarding (*or *cut-through*) devices, latency values are based on the last bit out the transmitting port and last bit in the receiving port. This result type is available in all the latency tests, including Jumbo, Latency, Latency Distribution, Latency Over Time, and Latency SnapShot. For testing *store-and-forward* devices, latency values may also be based on the first bit out the transmitting port and first bit in the receiving port. This result type is available in the Jumbo and Latency tests. For these tests, you select which result type you wish to view: cut-through or store and forward.

For further information, see *"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349* and *"Interpreting Latency Test Results" on page 352*.

Class of Service testing

Because SmartFlow maps into Layer 2, Layer 3, and Layer 4 of the OSI model, you can test a Class of Service (CoS) device for priority routing performance at any of these levels.

Tip: The SmartFlow online Help contains more information about how SmartFlow works and the calculations it uses. Information includes the calculation for the interburst gap plus how sequence tracking and data integrity works.

Test Flow Chart

Figure 2-5 illustrates the general methodology for each SmartFlow test. This chart roughly corresponds to the test phases detailed in the section.

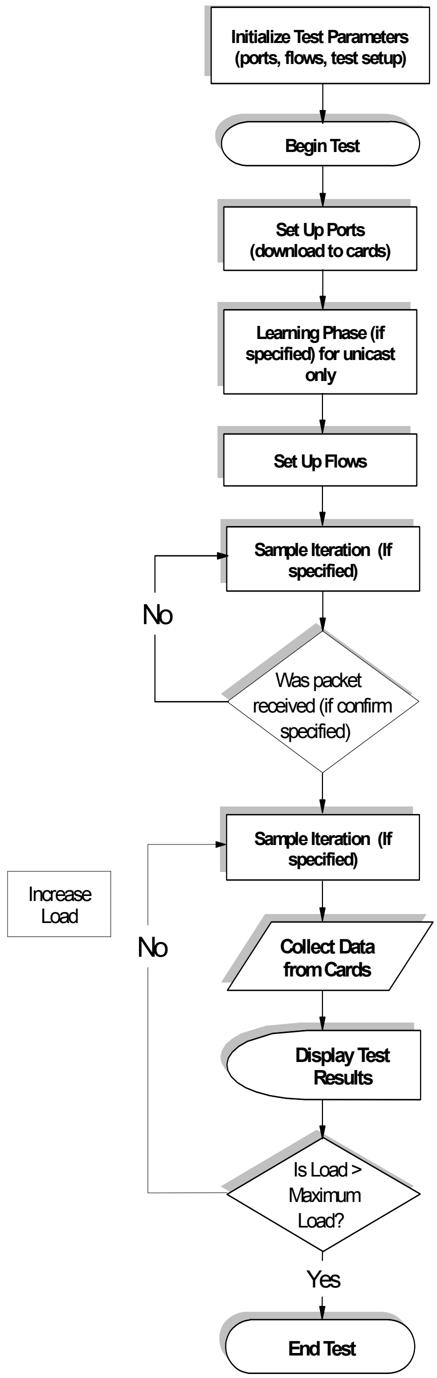

Figure 2-5. Test Flow Chart

Test Phases

Once you launch a test, SmartFlow runs through the following phases:

1. SmartFlow analyzes the setup data entered before running the test and builds a tracking system. It validates that flow IP addresses are legal according to the port and network setup. If they are not, an error message alerts you so that you can correct it before running the test.
2. SmartFlow checks for valid port setups and makes sure all test ports are ready to transmit.
3. If specified, Ethernet ports send out ARP requests or learning packets (for unicast traffic only, not multicast traffic), depending on how you configured the ports in SmartFlow. If learning fails, a message box displays the port for which ARP requests or learning failed. You will have the option to abort, ignore, or retry sending out these packets for the number of times specified on the *Test Setup>Learning* tab. See *"Address Resolution" on page 53* for more information.
4. SmartFlow downloads the specified test setups and card setups to the cards in the test.
5. If ports are enabled for multicast traffic, the multicast group receiver ports send IGMP Join requests. For more information about multicast traffic, see *Chapter 19, "IP Multicast Testing."*
6. If you selected to run a sample test iteration, SmartFlow runs a sample test iteration and then discards the results of this iteration. If you selected to confirm packet receipt, SmartFlow checks for receipt of at least one packet at each receiving port per flow.

The test phase of the test starts here.

7. If enabled for an Ethernet port, address resolution begins here. Ports sending IPv4 traffic send out ARP requests for the transmitting hosts (IP addresses) to the DUT. Ports sending IPv6 traffic exchange a neighbor solicitation and neighbor advertisement.
8. Test ports begin transmitting traffic for the duration specified and at the initial load specified. They send traffic in a burst consisting of the number of packets specified.
9. SmartFlow adjusts the traffic load based on the starting percentage of wire rate, the percentage to step up the rate, and the maximum rate or based on a custom load sequence. For a Throughput test in binary search mode, it adjusts the next traffic rate according to the binary search results, the backoff rate, and the resolution rate.
10. Ports stop transmitting traffic when the transmission rate reaches the specified maximum. For a Throughput test in binary search mode, ports stop transmitting when the required adjustment percentage (up or down) in the transmission rate is smaller than the specified *Resolution* value.
11. If any ports are enabled for multicast traffic, the multicast group receiver ports send IGMP leave requests at the end of the test if you selected this option.
12. SmartFlow gathers data from the ports and counters are updated.

13 SmartFlow automatically displays test results in the main window in chart and table format. It logs test results into files that you can view outside SmartFlow and print with standard utilities.

For information about performing the learning phase independent of running a test, see *"Running the Learning Phase" on page 296*.

How SmartFlow Computes Default IP Addresses for Flows and Ports

SmartFlow computes default IP addresses (IPv4 and IPv6) for IP ports and:

- Uses the same network (IPv4) or on-link (IPv6) as the gateway for the port.
- Bases the least significant byte (IPv4) or 16-bit integer (IPv6) on (but not equal to) the port number. (The least significant byte or 16-bit integer is based on the relative position of the port on the *Cards* tab.)
- Is unique from other IP addresses in that network or on-link.

Note: SmartFlow automatically converts default IPv4 addresses on the *IPv4 Networks* tab to default IPv6 addresses on the *IPv6* tab. It does this by prepending 96 zero bits to the IPv4 address. However, if you modify the default IPv4 address for a port, the corresponding IPv6 address does *not* change.

For information about how SmartFlow increments IP addresses if you chose to vary fields within a flow, see *"How SmartFlow Varies IP Addresses" on page 195*.

Sequence Tracking

The Frame Loss and Jumbo tests provide sequence tracking statistics, such as the number of packets received that were in sequence and out of sequence.

SmartFlow tracks the order of packets received through a sub-field in the proprietary signature field, which contains a frame sequence number. This sequence number starts at zero for each flow. It is incremented by one each time a packet is transmitted from the port. The sequence number is limited to 16 bits and therefore wraps every 65,536 packets.

What is an Out-of-Sequence Packet?

SmartFlow determines whether or not a packet is out of sequence by looking at the sequence number of the previously transmitted packet. If a packet sequence number is not the *expected* one (previous packet sequence number + one), the packet is considered out of sequence. The out-of-sequence counter is incremented by one.

Example: No packet loss occurred, and the DUT reordered only one packet (packet 4).

Assume this sequence of test frames was sent by the SmartBits port and received by the DUT:

1, 2, 3, 4, 5, 6, 7, 8, 9, 10

This is what occurred:

Received	1,	2,	3,	5,	6,	4,	7,	8,	9,	10
Expected	1	2	3	4	6	7	5	8	9	10
Sequence	*i*	*i*	*i*	*o*	*i*	*o*	*o*	*i*	*i*	*i*

Number of packets in sequence (i): **7**
Number of packets out of sequence (o): **3**

How Packet Loss Relates to Packet Sequencing

Packet loss affects the number of packets that are received in sequence. Therefore, tests that track packet sequencing are most effective when there is no packet loss or duplication. It is recommended that you run a Throughput test first, to determine the rate at which the DUT does not lose packets.

Example: Only one packet (packet 4) is not received. The next packet (packet 5) is considered out of sequence.

Assume this sequence of test frames was sent by the SmartBits port and received by the DUT:

1, 2, 3, 4, 5, 6, 7, 8, 9, 10

This is what occurred:

Received	1,	2,	3,	5,	6,	7,	8,	9,	10
Expected	1	2	3	4	6	7	8	9	10
Sequence	*i*	*i*	*i*	*o*	*i*	*i*	*i*	*i*	*i*

Number of packets in sequence (i): **8**
Number of packets out of sequence (o): **1**

Packet 4 is a lost (dropped) packet, but it is counted as an out-of-sequence packet.

Address Resolution

Address resolution, when enabled in SmartFlow, varies both in name and how it is done according to whether you are sending IPv4 or IPv6 traffic. If the *Addr Resolution* field is enabled on the *Cards* tab, ports send out either:

- ARP requests and/or learning packets (IPv4) or
- Neighbor solicitation/advertisement packets (IPv6).

ARP Requests/Learning Packets (IPv4)

ARP requests and learning packets enable the DUT and SmartBits ports to learn one another's addresses, so that they can send and receive IPv4 traffic between them. SmartFlow sends either ARP requests or learning packets during the *learning* phase, which takes place before the test phase of a test. This is done to populate the DUT's forwarding table or ARP cache table before the test phase, so that the test is not affected by these packets and they are not counted in test results.

Both learning and sending ARP requests consist of two phases:

1. Source ports send ARP or learning frames to the DUT to build the DUT's forwarding table or ARP cache table. Source ports also send these frames for each cyclic flow (if any) associated with the port.
2. *If you choose to run a sample test iteration,* for each flow the source ports send actual test frames configured to destination ports to verify the DUT's forwarding table. For example, if you are using 5 ports in a full mesh pattern, SmartFlow sends 10 packets for each source and destination of each flow that exists, yielding 250 test frames sent to verify the forwarding table.

For IPv4, SmartFlow can send out either ARPs or learning packets, depending on how you set up your test. If ARPs are enabled (*Addr Resolution* field is selected on the *Cards* tab), SmartFlow sends ARPs during the learning phase.

Notes:
- POS, WAN, and ATM cards do not require learning and do not need to build a forwarding table.
- You can set the frequency of when ARPs are sent out on the *Learning* tab in test setup. See *"Setting Up Test Parameters" on page 216* for more information.
- If the port has VLAN tags enabled (*IPv4 Networks* tab), ARP requests or learning packets will also contain VLAN tags.
- If you want to pre-populate the DUT's forwarding tables before actual test traffic is sent, run a sample test iteration. This is also known as "fast-path learning." See *"Sample Iteration Tab" on page 222* for more information.

The following paragraphs summarize the characteristics of ARP requests and learning packets.

ARP (Address Resolution Protocol) requests

ARP requests are used at Layer 3 when ports are on either the same or different subnets to map IP addresses to physical MAC addresses. ARPs are sent to enable the DUT to associate the MAC and IP addresses for each SmartFlow flow involved in the test as well as for the SmartBits port to learn the DUT's MAC address. You specify sending ARPs by selecting the *Addr Resolution* field on the *Cards* tab.

How ARPs work

During the learning phase, each port that has ARP enabled sends out ARP requests, which are broadcast onto the network with the router's destination IP address. The router responds to each port with its MAC address and updates the values in its address table with the IP address of each flow. SmartFlow updates the destination MAC address to that of the device.

If the destination IP address is on a different network than the source IP address, then the ARP request will be sent to the gateway (with a MAC address of all F's in the ARP request). If the destination IP address is on the same network as the source IP address, then the ARP request will be sent to the destination IP address.

If source and destination ports are on the same network or subnet, the destination port ARP replies to the destination port. If the destination port is on a different network or subnet, the gateway ARP replies.

Learning packets

Learning packets initialize device forwarding tables and are used at Layer 2 as a way for the DUT to register the MAC addresses of each flow involved in the test. Learning packets are ordinary IP packets containing the MAC addresses with which you want to populate the DUT's forwarding table. You specify sending learning packets at the *Learning* tab of the *Test Setup* tab. If ARPs are enabled, SmartFlow sends only ARPs.

How Learning Packets work

During the learning phase, all flows send learning packets so that the device can store the flow's source and MAC addresses in its address table.

Note: The MAC address frame format is changed during the learning phase.

The destination MAC address in the learning frame reflects the destination address of the corresponding test frame.

If the *Learning options* field on the *Test Setup>Learning* tab is set to any value except **NONE**, each port (except POS, WAN, and ATM ports) involved in the test sends learning packets or ARPs to devices.

Note: POS, WAN, and ATM ports do not require learning and do not need to build a forwarding table.

Neighbor Discovery (IPv6) in SmartFlow

IPv6 does not use ARP. Instead, it uses *neighbor discovery* procedures to associate an IP address with the Ethernet MAC address of its node. SmartFlow currently supports the neighbor solicitation/advertisement part of IPv6 neighbor discovery. SmartFlow performs the following steps for *each* SmartBits port that sends or receives IPv6 traffic. (See *Figure 2-6*.)

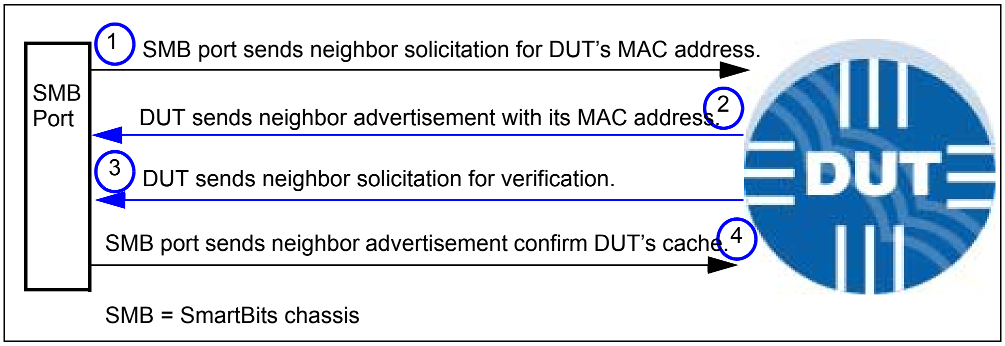

Figure 2-6. Neighbor Discovery in SmartFlow

1 The SmartBits port acts as a node and sends a *neighbor solicitation* packet to the DUT to request the link-layer (MAC) address. It waits for a neighbor advertisement packet from the DUT before sending any test packets.

2 The DUT sends a *neighbor advertisement* packet containing its MAC address to the SmartBits port.

3 To verify the transmitting port's IP/MAC address, the DUT sends a *neighbor solicitation* packet to the port.

  a The default time that SmartFlow waits for a DUT to send a neighbor solicitation is 3 seconds. To change the wait time for a neighbor solicitation from the DUT, add this statement to the end of the SmartFlow.ini file:

```
[Flow Core]
NeighborDiscoveryWaitTime=3000
```

  Change 3000 to a larger value. Valid values: 0 to 4294967295 (milliseconds).

  b Most DUTs do address verification at this point. However, some DUTs postpone verification until the DUT receives a test packet destined for that port. For this type of DUT, SmartFlow can send out an IP packet from the transmitting port prior to the start of the test in order to complete the address verification. To do this, add this statement to the end of the SmartFlow.ini file:

```
[Flow Core]
ProbeDuringNeighborDiscovery=1
```

4 The SmartBits port sends out a neighbor advertisement packet to the DUT to confirm the DUT's address cache.

Stream Maximums per Card and Learning/Address Resolution

Each card/module has a maximum number of hardware streams. Keep in mind that each address resolution/learning request and reply adds to this number of streams when you:

- Enable the **Address Resolution** field (on the *Cards* tab).
- Select the extent to which you want learning to occur [*Full(Tx+Rx), Partial(Rx), or None*] in the *Learning options* field on the *Test Setup>Learning* tab during test setup.

For example, if you have three source IP addresses (hosts) and two destination IP addresses in a non-cyclic flows, you could have five hardware streams depending on how your flows are set up. Therefore, particularly if you plan to have a large number of flows per port, it is recommended that you estimate the number of actual hardware streams for all source and destination IP addresses in non-cyclic flows involved in the test. Make sure their number does not exceed the card's maximum.

For information about the maximum number of streams per card, see *"Maximum Number of Flows per Port" on page 45*.

How Configurations Affect the Number of Streams

The resulting number of streams also depends in part on your configuration. For example, if you have many transmitting ports all sending to only one receiving port, that receiving port could exceed the card limit on the number of streams.

Non-cyclic SmartFlows that have variable fields use more hardware streams than those with non-variable fields. For more information about SmartFlows and streams, see *"How a Stream Generates a Flow" on page 32*.

From the *Test Setup>Learning* tab, if you select *Full* **(**learning packets for both transmitting and receiving ports), more streams are required than if you select *Partial (Rx)* during the learning phase.

Chapter 3

Install and Connect

This chapter provides information on the SmartFlow installation requirements, the SmartFlow installation procedure, and procedures for connecting your PC, chassis, and DUT/SUT.

For detailed information about cables and the physical connections between the chassis and your PC, see the installation guide for your chassis.

In this chapter...

Summary of Steps to Set up and Run a Test

Once you have installed SmartFlow, follow these steps to use the application:

1. Prepare the DUT.
2. Physically connect the SmartBits chassis to either your PC or to the network to which your PC is connected.
3. *Optional:* If you are using multiple SmartBits chassis, synchronize the chassis. (Refer to the *Installation Guide* for your chassis for more information.)
4. Physically connect cables from the SmartBits ports to be used in the test to the DUT.
5. Start the SmartFlow application.
6. Select **Setup>Chassis Connections** from the main menu to specify the SmartBits chassis to connect to and connection type.
7. Select **Action>Connect** to connect SmartFlow to the SmartBits chassis.
8. Using the *Cards* tab, configure the ports that you want to include in the test.

   Note: Once you connect to a multi-user chassis and reserve the SmartBits ports that you want to use, no other users can use those ports until you release them. To maximize card availability for other users, you can postpone this step until you are ready to start running the test. If you do not reserve the ports, SmartFlow automatically reserves them when the test is run.
9. Define and add the networks and VLANs that you want associated with each port at the *IPv4 Networks* and/or *IPv6* tab.
10. Plan what type of flows and groups that you want to create.
11. Use the **Group Wizard** on the *Groups* tab to set up flows and groups.
12. Specify test parameters at the *Test Setup* tab.
13. Click the appropriate test **Launch** button on the shortcut bar to start the test.
14. View results in the *Results* window of any results tab, such as the *Chart*, *Summary*, or *Detail* tabs. Click the **Results** button on the shortcut bar if results do not appear.

SmartBits Hardware Requirements

Table 3-1 through *Table 3-4* list the SmartBits hardware requirements:

- *Table 3-1* – lists supported SmartBits chassis
- *Table 3-2* – lists supported cards and modules
- *Table 3-3* – lists supported TeraMetrics modules
- *Table 3-4* – lists supported TeraMetrics-based modules.

SmartFlow operation requires at least one of the supported chassis types and two or more of the supported SmartCards or modules.

Table 3-1. Supported SmartBits Chassis

Chassis
SmartBits 200
SmartBits 2000
SmartBits 600
SmartBits 600B
SmartBits 6000B
SmartBits 6000C

Cards and modules

At least two ports on any of the following cards or modules in any combination must be used. For current firmware versions of SmartBits cards/modules chassis, see the *Release Notes* (`readme.pdf`) included on the product CD.

Table 3-2. Supported SmartCards and Modules

Topology	SmartCard / Module	Description
LAN	ML-7710	10/100Base-TX Ethernet, 1-port, SmartMetrics SmartCard
	ML-7711 ML-7711s	100Base-FX Ethernet, 1-port, multi-mode, 1300nm, SmartMetrics SmartCard ML-7711s 100Base-FX Ethernet, 1-port, single mode, 1310nm**,** SmartMetrics SmartCard
	ML-5710A	10Base-T Ethernet/USB, 2-port, SmartMetrics SmartCard
	LAN-3101A/B	10/100Base-T Ethernet, 6-port, SmartMetrics module

Table 3-2. Supported SmartCards and Modules (continued)

Topology	SmartCard / Module	Description
LAN	LAN-3102A	10/100Base-T Ethernet, 2-port, SmartMetrics module
	LAN-3111A LAN-3111As	100Base-FX Ethernet, 6-port, multi-mode, 1300nm, SmartMetrics module 100Base-FX Ethernet, 6-port, single mode, 1310nm, SmartMetrics module
	LAN-3201B/C	1000Base-X Ethernet, GBIC, 1-port, SmartMetrics module
	LAN-3300A LAN-3301A	10/100/1000Base-T Ethernet, Copper, 2-port, SmartMetrics module 10/100/1000Base-T Ethernet, Copper, 2-port, TeraMetrics module
	LAN-3302A	10/100Base-TX Ethernet, Copper, 2-port, TeraMetrics module
	LAN-3306A	10/100Base-TX Ethernet, Copper, 4-port, TeraMetrics XD module
	LAN-3310A LAN-3311A	1000Base-X Ethernet, GBIC, 2-port, SmartMetrics module 1000Base-X Ethernet, GBIC, 2-port, TeraMetrics module
	LAN-3320A	10/100/1000 Mbps and Gigabit Ethernet Fiber, 2-port, SmartMetrics XD module
	LAN-3321A LAN-3324A LAN-3325A LAN-3327A	10/100/1000Base Ethernet and Gigabit Fiber, 2-port, TeraMetrics XD module 10/100/1000 Mbps and Gigabit Ethernet Fiber, 4-port, SmartMetrics XD module 10/100/1000 Mbps and Gigabit Ethernet Fiber, 4-port, TeraMetrics XD module 10/100/1000 Mbps and Gigabit Ethernet Fiber, 1-port, TeraMetrics XD module
LAN/WAN	XFP-3730A XFP-3731A	10GBase Ethernet, XFP MSA, 1-port,1-slot, SmartMetrics module 10GBase Ethernet, XFP MSA, 1-port, 1-slot, TeraMetrics module
	XLW-3720A XLW-3721A	10GBase Ethernet, XENPAK MSA, 1-port, 2-slot, SmartMetrics module 10GBase Ethernet, XENPAK MSA, 1-port, 2-slot, TeraMetrics module
POS	POS-3500B POS-3500B(s)	OC-3c/OC-12c (STM-1c/STM-4c), 1-port, multi-mode, 1300nm, SmartMetrics module OC-3c/OC-12c (STM-1c/STM-4c), 1-port, single mode, 1310nm, SmartMetrics module
	POS-3502A POS-3502A(s)	OC-3c (STM-1c), 1-port, multi-mode, 1300nm, SmartMetrics module OC-3c (STM-1c), 1-port, single mode, 1310nm, SmartMetrics module
	POS-3504As POS-3504Ar	OC-48c (STM-16c), 1-port, single mode, 1310nm, SmartMetrics module OC-48c (STM-16c), 1-port, single mode, 1550nm, SmartMetrics module

Table 3-2. Supported SmartCards and Modules (continued)

Topology	SmartCard / Module	Description
	POS-3505As POS-3505AR	OC-48c (STM-16c), 1-port, single mode, 1310nm, TeraMetrics module OC-48c (STM-16c), 1-port, single mode, 1550nm, TeraMetrics module
	POS-3510A POS-3510As	OC-3c/OC-12c (STM-1c/STM-4c), 2-port, multi-mode, 1300nm, SmartMetrics module OC-3c/OC-12c (STM-1c/STM-4c), 2-port, single mode, 1310nm, SmartMetrics module
	POS-3511A POS-3511As	OC-3c/OC-12c (STM-1c/STM-4c), 2-port, multi-mode, 1300nm, TeraMetrics module OC-3c/OC-12c (STM-1c/STM-4c), 2-port, single mode, 1310nm, TeraMetrics module
	POS-3518As POS-3518AR	OC-192c (STM-64c), 1-port, 2-slot, single mode, 1310nm, SmartMetrics module OC-192c (STM-64c), 1-port, 2-slot, single mode, 1550nm, SmartMetrics module
	POS-3519As POS-3519AR	OC-192c (STM-64c), 1-port, 2-slot, single mode, 1310nm, TeraMetrics module OC-192c (STM-64c), 1-port, 2-slot, single mode, 1550nm, TeraMetrics module
WAN	WN-3445A	DS3, Frame Relay/PPP, 1-port, SmartMetrics SmartCard
ATM	ATM-3451A ATM-3451As	OC-3c (STM-1c), 2-port, multi-mode, 1300nm, TeraMetrics module OC-3c (STM-1c), 2-port, single mode, 1310nm, TeraMetrics module
	ATM-3453A ATM-3453As	OC-3c/OC-12c (STM-1c/STM-4c), 2-port, multi-mode, 1300nm, TeraMetrics module OC-3c/OC-12c (STM-1c/STM-4c), 2-port, single mode, 1310nm, TeraMetrics module

Notes:

- For information on the latest firmware levels for the supported SmartCards and modules, refer to the SmartFlow Release Notes shipped with the software.
- To check the firmware of the SmartCards/modules installed in your chassis, select **Help>About** within the SmartFlow application.
- Link failures that cause fluctuation on the link statue could cause erroneous results on 10Gig cards.

Supported TeraMetrics and TeraMetrics-based Modules

Table 3-3 lists the supported TeraMetrics modules. *Table 3-4* lists the supported TeraMetrics-based modules. All these modules are also included in the complete list in *Table 3-2 on page 59*.

Table 3-3. Supported TeraMetrics Modules

Topology	TeraMetrics Module	Description
ATM	ATM-3451A ATM-3451As	ATM OC-3c (STM-1c), 2-port, multi-mode, 1300nm, TeraMetrics module ATM OC-3c (STM-1c), 2-port, single mode, 1310nm, TeraMetrics module
	ATM-3453A ATM-3453As	ATM OC-3c/OC-12c (STM-1c/STM-4c), 2-port, multi-mode, 1300nm, TeraMetrics module ATM OC-3c/OC-12c (STM-1c/STM-4c), 2-port, single mode, TeraMetrics module
LAN	LAN-3301A	10/100/1000Base-T Ethernet, Copper, 2-port, TeraMetrics module
	LAN-3302A	10/100Base-TX Ethernet, Copper, 2-port, TeraMetrics module
	LAN-3306A	10/100Base-TX Ethernet, Copper, 4-port, TeraMetrics XD module
	LAN-3321A	10/100/1000 Mbps and Gigabit Ethernet Fiber, 2-port, TeraMetrics XD module
	LAN-3325A	10/100/1000 Mbps and Gigabit Ethernet Fiber, 4-port, TeraMetrics XD module
	LAN-3327A	10/100/1000 Mbps and Gigabit Ethernet Fiber, 1-port, TeraMetrics XD module
LAN/WAN	XLW-3721A	10GBase Ethernet, XENPAK MSA, 1-port, 2-slot, TeraMetrics module
	XFP-3731A	10GBase Ethernet, XFP MSA, 1-port, 1-slot, TeraMetrics module
POS	POS-3505As POS-3505Ar	POS OC-48c (STM-16c), 1-port, single mode, 1310nm, TeraMetrics module POS OC-48c (STM-16c), 1-port, single mode, 1550nm, SmartMetrics module
	POS-3511A POS-3511As	POS OC-3c/OC-12c (STM-1c/STM-4c), 2-port, multi-mode, 1300nm, TeraMetrics module POS OC-3c/OC-12c (STM-1c/STM-4c), 2-port, single mode, 1310nm, TeraMetrics module
	POS-3519As POS-3519Ar	POS OC-192c (STM-64c), 1-port, 2-slot, single mode, 1310nm, TeraMetrics module POS OC-192c (STM-64c), 1-port, 2-slot, single mode, 1550nm, TeraMetrics module

Table 3-4. Supported TeraMetrics-based Modules

Topology	TeraMetrics-based Module	Description
LAN	LAN-3300A	10/100/1000Base-T Ethernet Copper, 2-port, SmartMetrics module
	LAN-3320A	10/100/1000 Mbps and Gigabit Ethernet Fiber, 2-port, SmartMetrics XD module
	LAN-3324A	10/100/1000 Mbps and Gigabit Ethernet Fiber, 4-port, SmartMetrics XD module
LAN/WAN	XLW-3720A	10GBase Ethernet, XENPAK MSA, 1-port, 2-slot, SmartMetrics module
	XFP-3730A	10GBase Ethernet, XFP MSA, 1-port, 1-slot, SmartMetrics module
POS	POS-3504As POS-3504AR	POS OC-48c (STM-16c), 1-port, single mode, 1310nm, SmartMetrics module POS OC-48c (STM-16c), 1-port, single mode, 1550nm, SmartMetrics module
	POS-3510A POS-3510As	POS OC-3c/OC-12c (STM-1c/STM-4c), 2-port, multi-mode, 1300nm, SmartMetrics module POS OC-3c/OC-12c (STM-1c/STM-4c), 2-port, single mode, 1310nm, SmartMetrics module
	POS-3518As POS-3518AR	POS OC-192c (STM-64c), 1-port, 2-slot, single mode, 1310nm, SmartMetrics module POS OC-192c (STM-64c), 1-port, 2-slot, single mode, 1550nm, SmartMetrics module

PC Requirements

- An IBM or compatible PC with an available COM port or an Ethernet interface card
- Pentium II processor 266 Mhz or higher
- One of the following Windows operating systems:
  - Windows 2000, Professional Edition, Service Pack 4
  - Windows NT 4.0, Service Pack 6a
  - Windows XP Professional Edition, Service Pack 1, 1a, 2
- 256 MB of RAM for Windows 2000/NT/XP
- 1.5 GB free hard disk space (for installation)
  If you plan to install Internet Explorer in order to run the HTML Help file, you need at least an additional 30 MB.
- SVGA monitor with 800 x 600 or better resolution
- *For large-scale testing:*
  - 1.5 GHz processor
  - 512 MB of RAM

  Large scale tests may include one or more of the following:
    - Multiple chassis tests with multiple cards in each chassis
    - Several hundred or more test iterations
    - Several thousand or more flows.
- *Optional:* For chassis synchronization, either a synchronization cable or two GPS receivers and antennas, plus an additional SmartBits 200 or 2000 chassis. (Refer to the *Using GPS/CDMA with SmartBits* document that accompanies the GPS unit for more information.)

Notes:
- The maximum number of flows, test iterations, networks, and VLANs that you can use is subject to the amount of memory on the PC that is running SmartFlow.

Third-party Software

- Microsoft Excel 2000 or later
- MS Internet Explorer 5.00 or later
- MS Outlook 98 or later

Installing SmartFlow

SmartFlow can be used with Windows 2000/NT/XP systems. (See *"PC Requirements" on page 64* for the specific operating systems that are supported.) SmartFlow ships on a CD-ROM and is also available from the Spirent Communications Customer Service Center website at http.support@spirentcom.com.

To install SmartFlow:

1 *To install SmartFlow from a CD*, insert the CD into your PC CD-ROM drive. The SmartFlow installer program launches automatically.

To install SmartFlow from the Spirent Communications Customer Service Center website:

- a Login to the website and select **Software Downloads>All Software** from the menu bar.
- b Under the *SmartBits* heading, click **Applications** and then click on the SmartFlow listing.
- c Click the **Download** button on the *Software Downloads* SmartFlow web page to copy the `.exe` file to your PC, then click on the `.exe` file to launch the SmartFlow installer program.

2 Follow the instructions on the installer screens to install <Product Name> properly.

3 After installing SmartFlow, you need to install the key file that enables you to use SmartFlow with your particular SmartBits chassis. This key file is contained on the 3.5-inch disk that is shipped with the SmartFlow CD. Insert the disk into your PC and locate the key file, which appears in the format key.txt (FLOWKEY.txt). Copy this file into the SmartFlow program directory.

Notes:

- *If you are using a purchased version of* SmartFlow*:* If the key file is not present in the program directory (or if the key file does not match the chassis to which you are trying to connect), the application starts up in a 15-minute demo mode. After 15 minutes, the application disconnects from the chassis. You can restart the application for another 15 minutes as many times as you want.
- *If you are using a demo (10- or 30-day* trial*) version of* SmartFlow*:* Demo versions allow you to connect to any chassis and use the software for the extent of the trial period. Once the trial period is over, you are permanently locked out of the software.
- To update the firmware for the chassis or modules you plan to use for SmartFlow testing, follow the instructions in the Release Notes for the firmware. You can also use the Download Manager online help.
- To run SmartFlow with multiple chassis, please contact Spirent Communications Technical Support department for a multiple chassis key file. (See *"How to Contact Us" on page 14*.)

Connection Procedures

Before you can begin running tests, you must:

- Physically connect your PC to a SmartBits chassis using either an Ethernet or Serial connection. (To do this, see the *SmartBits 200/2000 Installation Guide* or *SmartBits 600x/6000x Installation Guide* that was shipped with your chassis.)
- Set up SmartBits connections using SmartFlow.
- Establish a connection to a SmartBits chassis using SmartFlow.
- Connect the SmartBits chassis to a DUT/SUT.

Synchronizing Chassis

If you are using multiple chassis, synchronize the chassis by using either:

- An expansion in and expansion out port of the chassis and a synchronization cable.
- A Global Positioning System (GPS) by connecting a GPS receiver to the SmartBits chassis and then synchronizing clocks between the various chassis using GPS. (For more information, refer to the documentation enclosed with your GPS equipment.)
- Code Division Multiple Access (CDMA) by connecting a CDMA receiver to the SmartBits chassis and then synchronizing clocks between the various chassis using CDMA. (For more information, refer to the documentation enclosed with your CDMA equipment.)

Important: When using a multiple chassis arrangement that contains SmartBits 600x/6000x chassis, you must use synchronization cables to connect the chassis together. In the master chassis, this synchronization cable connects to the out port only. In the last chassis in the multiple chassis arrangement, the synchronization cable must connect to the in port only. With all other chassis in the arrangement, cables must connect to both the in and out ports.

An appropriate CAT-5, Ethernet straight-through sync cable is shipped with every SmartBits chassis. If you do not have this cable or need more, you can substitute a CAT-5, straight-through (not crossover) cable that is 1m or less in length. The use of this cable is required in order for overall test results to be accurate (not just Latency test results, but all test results).

The synchronized clock and pulse signals passing through the connected chassis require careful sequential powering up. The first chassis that is booted up supplies the master clock to all of the other chassis.Wait 10 seconds and then boot the next chassis in the chain. Wait another 10 seconds and then boot the next chassis, and so on.

For detailed information on synchronization and multiple chassis operation, refer to the *SmartBits 600x/6000x Installation Guide* or to Application Note #18, *Connecting & Synchronizing Multiple SmartBits Chassis*.

Preparing the DUT

1. Identify the ports that you want to test.
2. Ensure that each port to be tested is configured on the DUT.

Connecting a SmartBits Chassis to a PC

To connect the SmartBits chassis to either your PC or to the network to which your PC is connected, use one of these chassis ports labeled:

- 10/100Base-T for an Ethernet connection on a SmartBits 600x/6000x
  (For SmartBits 600x/6000x, use *only* Ethernet connections.)
- 10Base-T for an Ethernet connection on the SmartBits 200/2000
- CONSOLE for a serial connection on the SmartBits 200/2000. (On the SmartBits 600x/6000x, use *only* the serial connection for management functions such as setting the SmartBits chassis MAC and IP addresses.)

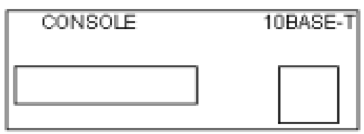

SmartBits chassis ports to use to connect to a PC

For detailed information, see your SmartBits chassis installation manual.

For information about the correct method to use to synchronize chassis, see the application note *Connecting Multiple SmartBits Chassis* which can be downloaded from the Support area of the Spirent Communications website www.spirentcom.com.

Starting SmartFlow

Before starting SmartFlow, make sure that the SmartBits chassis is connected to your PC and is powered on.

To start SmartFlow, follow these steps:

1 Either:

- On the Windows taskbar, click the **Start** button, and point to **Programs>SmartBits Applications>SmartFlow**, and then click on **SmartFlow**. (This path can vary if you did not use the default folder during installation.)
- Double-click on the **SmartFlow** icon from your desktop.

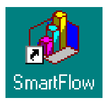

If you have already selected **Setup>Connect on Startup** in an earlier session, SmartFlow immediately attempts to establish a connection with the SmartBits chassis.

2 Click the **Connect** button on the toolbar or select **Action>Connect** to connect using the connections selected.

If the connection is successful, the connection light on the lower right portion of the main window turns green and the number of connections is displayed.

Connecting SmartFlow to SmartBits Chassis

Before you can select and configure test ports, you should first:

1 Set up the connection(s) to the chassis containing the ports that you plan to use.

Use the *Setup SmartBits Connections* dialog to define chassis to which you want to connect.

2 Connect SmartFlow to the SmartBits chassis. (See *“Connecting to Chassis in SmartFlow” on page 71* for more information.)

Note: It is recommended that you connect to the chassis *before* you set up the ports, networks, and flows so that flows and ports belong to the correct networks.

Setting Up Chassis Connections in SmartFlow

Once you have set up the connections, you can manually or automatically connect to active chassis connections whenever you start SmartFlow. You set up connections in the *Setup SmartBits Connections* dialog box (*Figure 3-1*).

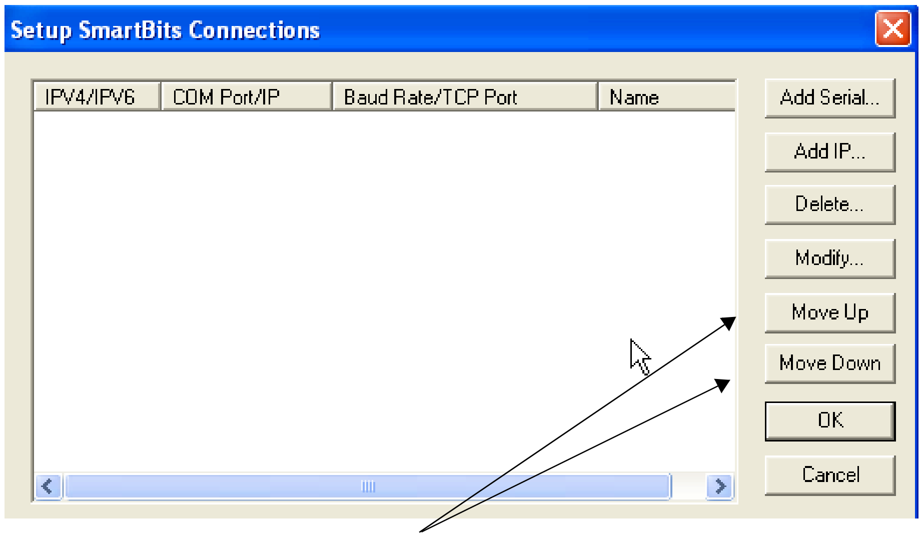

These buttons change the order in which connections are listed in the dialog box. The order here determines the order in which chassis and ports are displayed on the *Cards* tab.

Figure 3-1. Setup SmartBits Connections Dialog Box

SmartFlow allows you to connect to multiple chassis and multiple types of chassis, such as SmartBits 200 and SmartBits 6000x, utilizing ports from any combination of SmartBits chassis. You can set up and connect to a maximum of 16 SmartBits chassis connections over IPv4 or IPv6.

Tip: When setting up a connection, use the *Name* field to help distinguish between chassis in SmartFlow dialog boxes. If you do not enter a name, SmartFlow enters a default name after you click the **OK** button.

If you define a chassis connection in the *Setup SmartBits Connections* dialog box, but no longer want to connect to it, you can simply deactivate the connection instead of deleting it by clearing the checkbox beside the IP address or COM port.

To add, modify, or delete SmartBits chassis in SmartFlow:

1 From the main window, select **Setup>Chassis Connections** or **Action>Connect** from the menu bar.

2 In the *Setup SmartBits Connections* dialog box, set up an appropriate communications port.

- If you use the Ethernet port, click the **Add IP** button.

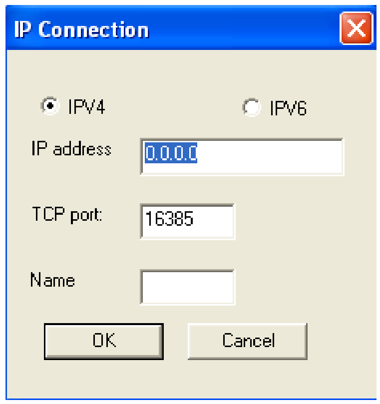

In the *IP Connection* dialog box, click the **IPv4** or **IPv6** button, depending on the IP version. Enter an IP address and TCP port number. (For SmartBits chassis, the TCP port number is always 16385.) You can also enter a description (such as QA-2000) for the chassis in the *Name* field. The name helps identify the chassis associated with the port when you connect to more than one chassis at once. Then click the **OK** button.

If you are not going through a router, make sure that the IP address used is located on the same subnet as your PC. For example, if the subnet mask is 255.255.255.0, the first three bytes of the address must be the same. For example:

If the PC IP address is:192.33.44. 5,
then the chassis IP address could be: 192.33.44. 6

Note: For information on IPv6 (including defining IPv6 information for a port, address formats, extension headers, and configuration limitations), refer to IPv6 index entries in this user guide.

- If you use the serial port, click the **Add Serial** button. In the *Serial Connection* dialog box, enter a COM port number and baud rate. Enter a description (such as QA-2000) for the chassis in the *Name* field. The name helps identify the chassis associated with the port when you connect to more than one chassis at once. Then click the **OK** button.

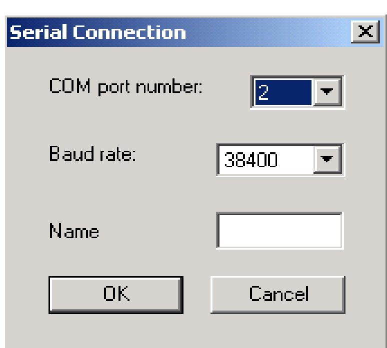

3. When you are ready to connect SmartFlow to the SmartBits chassis, click the **OK** button in the *Setup SmartBits Connections* dialog box.

   The connection light on the lower right portion of the main window should be green. The number of chassis connections is also displayed.

Note: If there is a problem with the connection (such as no power or an invalid ID), an error message appears that states:

`Error-2: Connect to SmartBits Failed on link: x.x.x.x`

Verify that the chassis is powered on, that a cable is not improperly connected to the Ethernet or serial port, and that the IP address for the chassis is correctly typed.

Connecting to Chassis in SmartFlow

You can connect to chassis from SmartFlow at either of these times:

- At SmartFlow startup:

  At the main window, select **Setup>Connect on Startup** from the menu bar. Whenever you start SmartFlow, it automatically attempts to establish a connection with all SmartBits chassis that are enabled in the *Setup SmartBits Connections* dialog box.

- Any time from the SmartFlow main window:

  Click the **Connect** button on the toolbar , select **Action>Connect** from the menu bar, or select **Setup>Chassis Connections** from the menu bar. Then click the **OK** button in the *Setup SmartBits Connections* dialog box.

During the connection process, you can cancel the process by clicking the **Cancel** button.

When connecting to *multiple* chassis, SmartFlow does the following:

1. SmartFlow attempts to establish a connection to the first chassis enabled in the *Setup SmartBits Connections* dialog box. It attempts this for 30 seconds, which is the default timeout period.
2. If it successfully connects, SmartFlow then attempts to establish a connection with the next enabled chassis. If it is able to connect, it continues to attempt connections with each of the other enabled chassis in the same manner.

   If it does *not* successfully connect to the first chassis within 30 seconds, SmartFlow displays an error message. Once you click the **OK** button, it attempts to establish a connection to the next enabled chassis for 30 seconds, and so on for each enabled chassis.

To change the default timeout period to connect to each chassis:

1 Open the `SmrtFlow.ini` file in a text editor.

2 Under the *General* section, modify the value for the parameter **ConnectionTime-Out(Seconds)** to the number of seconds that you want to allow for each connection attempt.

If You Have Problems Connecting

Table 3-5 shows some common problems that you might encounter when connecting to chassis from within SmartFlow.

Table 3-5. Connection Problems

Problem	Solution
Connection is not enabled (selected).	Go to the *Setup SmartBits Connections* dialog box and select the checkbox beside the connection.
Connection is enabled but still cannot connect.	Verify that the chassis is powered on and/or all enabled chassis in the *Setup SmartBits Connections* dialog box are powered on.
Card and/or chassis firmware is not the required version.	**1** Select **Help>About SmartFlow** from the menu bar to access the *About SmartFlow* dialog box. (See *"Viewing Version Information" on page 93* for more information.) **2** Compare the current firmware with the required firmware. **3** Update the card and/or chassis firmware, as required. Go to the Spirent Communications website to download the latest firmware at http://www.spirentcom.com.

Connecting SmartBits to the DUT

Connect cables from the ports on the SmartBits chassis to the DUT. Do this for all cards involved in the test. You can connect up to 16 SmartBits chassis to each other: either 16 SmartBits 600x/6000x chassis or 16 stacks, each consisting of either four daisy-chained SmartBits 2000s or a combination of daisy-chained SmartBits 2000 and SmartBits 10s.

Using appropriate cables, connect the SmartBits ports to the DUT in any order. Look at your device's link light to verify that the cable is correctly connected.

Note: If you are using a fiber card, make sure that the fiber cable is never bent to a radius of less than 1.5 inches.

SmartBits Connection Tips

- If you are having difficulty connecting to SmartBits via a serial connection from SmartFlow, verify that SmartBits is connected by way of a straight-through serial cable to your PC. SmartBits does not use a null-modem cable.

  Also ensure that the COM port selected in the **Setup>Chassis Connections** menu is the one that connects SmartFlow to the SmartBits chassis.
- If you are using a card with a fiber interface, label each fiber cable connecting SmartBits ports with the DUT ports. Reason: Fiber cables are not rigid and tend to hang in one group when not connected to a DUT, causing confusion when you are ready to connect them again.
- Disconnect all cables from the DUT that are not part of the test.
  These ports could be receiving traffic that make your frame loss or latency greater than it should be for the number of ports listed in your test.
- Make sure to terminate all cables at both the SmartBits end and at the DUT end. An unused port should not have a cable plugged in and dangling.
- Replace the ports' connector dust plugs and cable end caps when connectors are not in use.
- Use only a certified CAT-5, CAT-5E or CAT-6 cable in any SmartBits Ethernet port (both control port as well as test port).
- Use only yellow optical cable with Smartbits ports that support single mode optics; use only orange or gray cable with Smartbits ports that support multi-mode optics.
- Make sure that you fully insert all SmartBits cards and modules into the chassis slots and tighten all thumb screws.
- In crowded engineering lab conditions, avoid passing test cable connecting the SmartBits chassis and the DUT over or under another DUT or live equipment. If the test cables are very long, avoid looping them around any live equipment.

- SmartFlow does not allow hot-swapping of SmartBits cards.
  You must disconnect SmartFlow, power off the chassis, adjust the cards and connections, power on the chassis again, and then connect SmartFlow to the SmartBits chassis from the *Action* menu.
- When connecting multiple chassis, use only shielded cables supplied by Spirent Communications.
- If using sync cables, make sure they are 1m or less in length and properly connected.
- Avoid running a test or control cable over the AC power supply of an open DUT. This will most likely produce faulty measurement results.

For more information on setting up the hardware and your SmartBits chassis, see the installation manual appropriate to your SmartBits chassis.

Chapter 4

Basic Navigation

This chapter describes the main user interface features and options of SmartFlow, including the main windows, menu and toolbar commands, setup tabs, and dialog boxes.

In this chapter...

Main Window

The SmartFlow main window contains the following areas:

- Menu Bar
- Toolbar
- Shortcut Bar (launches tests and displays results if available)
- Results Window panel
- Test Setup tabs (The order of these tabs can be modified.)

Figure 4-1 shows the SmartFlow main window with the *SmartFlows* tab active.

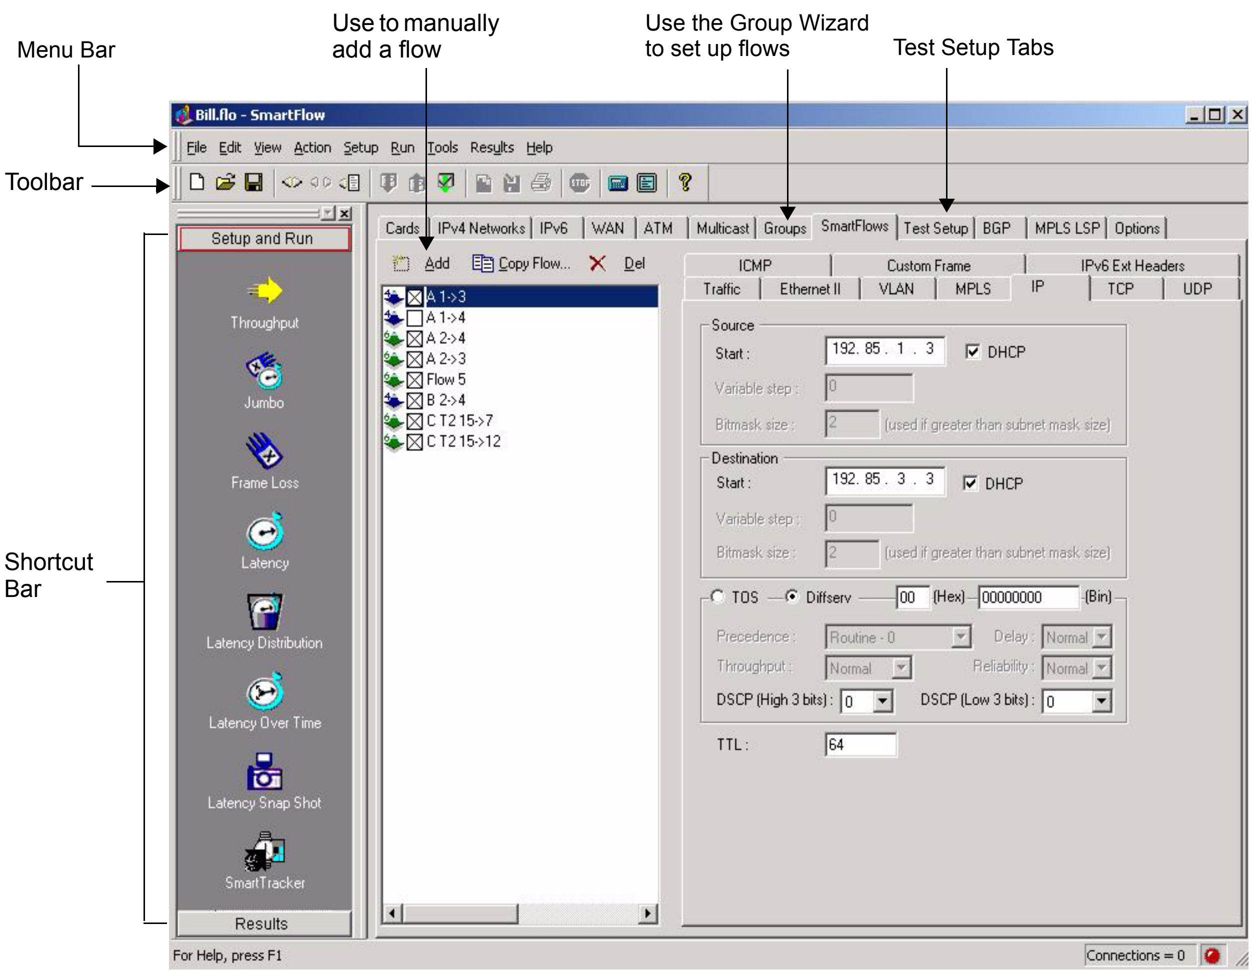

Figure 4-1. SmartFlow Main Window

Icons beside flows indicate whether the flow is an IPv4, IPv6, or custom frame flow.

Menus

The SmartFlow main window contains these menus: File, Edit, View, Action, Setup, Run, Tools, Results, and Help.

The SmartFlow online Help contains descriptions of every menu option.

Toolbar Buttons

The following table describes the buttons in the Toolbar at the top of the main window.

Button	Description
	Connects to a SmartBits chassis.
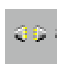	Disconnects from a SmartBits chassis.
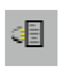	Sets up and changes SmartBits chassis connection settings.
	Initiates a DHCP request.
	Releases the DHCP IP addresses after a DHCP request is completed.
	Shows the DHCP connection status of each flow. (A white background indicates there is no DHCP configuration. A red background indicates the source and/or destination DHCP connection failed. A green background indicates both the source and destination DHCP connections resolved.)
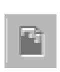	Copies the current test results to the Windows clipboard.
	Exports the current test results to a Windows file directory.

Button	Description
	Prints the current test results.
	Stops a test that is underway.
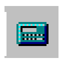	Opens the Flow Rate Calculator to calculate rates per flow for a port.
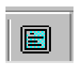	Enlarge the active *SmartFlow* tab to fill the full screen.
	Access SmartFlow online Help.

Shortcut Bar

The shortcut bar contains test launch buttons to start a test and two buttons at the extreme top and bottom to toggle between setup and results tabs. (You can also use the Ctrl + R keys to toggle between setup and results tabs.)

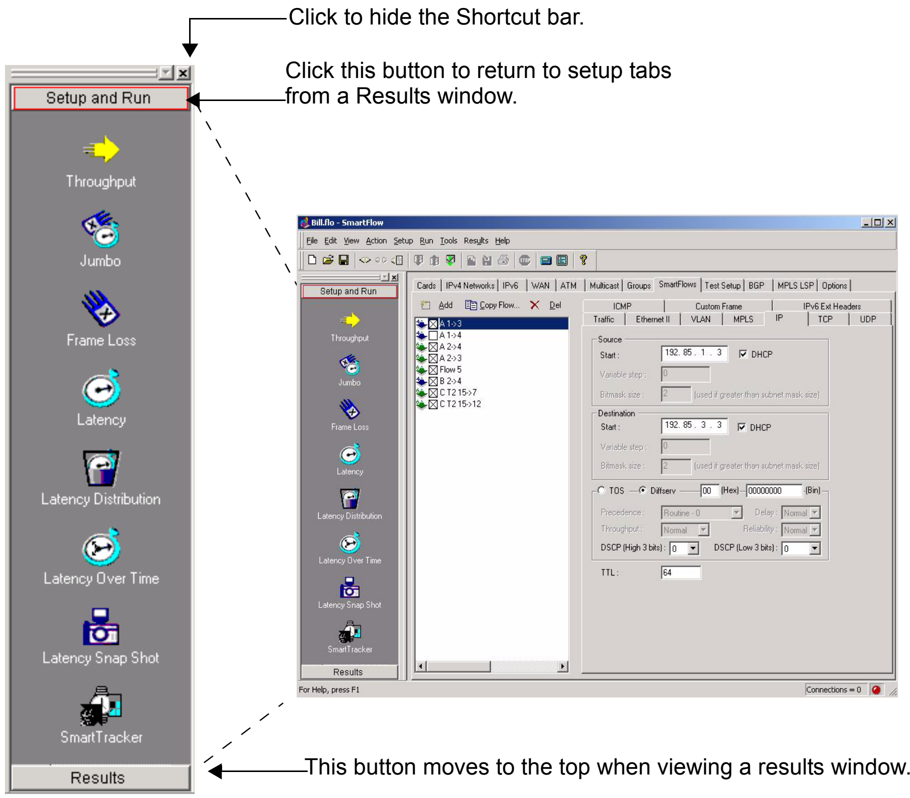

Once results are available, you can toggle between results and setup tabs by clicking the buttons shown at the top and bottom of the shortcut bar.

If the Results button is at the *bottom* of the shortcut bar, when you click on a test button that test launchs. If the Results button is at the *top* of the shortcut bar, if you click on a test button you see results for that test. (When test results are available, if you click on the Results button, it moves to the top of the shortcut bar.)

Note: You can arrange the test buttons on the shortcut bar in any order you want. Hold down the left mouse button and drag the icon to where you want it, then release the mouse button.

Showing/Hiding the Shortcut Bar

To hide the shortcut bar, click the small **x** button at the top of the bar. To show the shortcut bar, select **View** from the menu bar. Then click once to the left of the **Shortcut Bar** option on the pull-down menu.

Results Window

At the conclusion of each test, a Results window is displayed with the following tabs:

The SmartTracker test contains some additional tabs. See *Chapter 16, "SmartTracker Test"* for more information.

To access results, click the **Results** button on the shortcut bar.

Each results tab contains the following types of information:

- Chart

  Contains a graphical form of summary results for each group and/or flow. You can right-click to change the format between Manhattan and ribbon formats.

- Summary

  Contains high-level, summary information about each group and/or flow.

- Detail

  Contains detailed information about each group and/or flow. For more information, see *"Detail Tab" on page 81*.

- Stray Frames

  Contains information about each flow containing stray frames. These are frames that were received but were not destined (expected) for a particular port. The Latency Over Time test does not report stray frames results.

- Port Errors

  Contains information about any port errors that occurred during the test. For more information, see *"Port Errors Tab" on page 83*.

- Packet Rate

  Contains information about the number of transmitted and received packets per second for each port involved in the test. For more information, see *"Packet Rate Tab" on page 83*.

Since test result information varies from test to test, see the chapters that cover the tests.

Note: Chart and Summary tab information is based on intended, not offered (actual) loads. If you checked the **Offered (Actual) load** field, disregard results on these tabs.

To display both groups and flows in results, check the **Display flows in Summary views** option and **Display flows in Detail view** option on the *Options>Results* tab.

Information on each tab (except the *Port Errors* tab) varies by test. See the appropriate section and chapter for each test for field descriptions.

Scientific Notation in Chart Results

All chart test result values (except for the Frame Loss test) are presented in *scientific notation,* if the value is beyond six digits. (See *Figure 9-4 on page 322* for an example.) Scientific notation is used for very large numbers. The *e* and the number that follows it represents the exponent, which indicates how far the decimal point is to be moved if the number were represented in decimal notation.

For example, to represent 7.16428e8 in decimal notation, move the decimal point eight places to the right of the exponent which yields the number 716,428,000. The number 1.14881e6 is a greater value because it represents the decimal notation (integer) value of 1,148,810.0. The number 1.206e9 is larger than 7.16428e8 because once it is converted to decimal notation, it is 1206000000.0.

Detail Tab

When viewing detailed test results for the Jumbo, Throughput, Frame Loss, Latency, and Latency Distribution, you can change the frame size upon which to base the receive rate (*Rx L2 bps* field) calculation.

The *Adjustment for bps calculation* dialog (*Figure 4-2*) allows you to set a new non-payload (overhead) size by entering a larger number for the frame size. By default, the Layer 2 receive-rate calculation is based on the entire frame length. This includes the preamble (8 bytes), Layer 2 header (14 bytes), CRC (4 bytes), and IFG (12 bytes, the minimal size) for a total of 38 bytes.

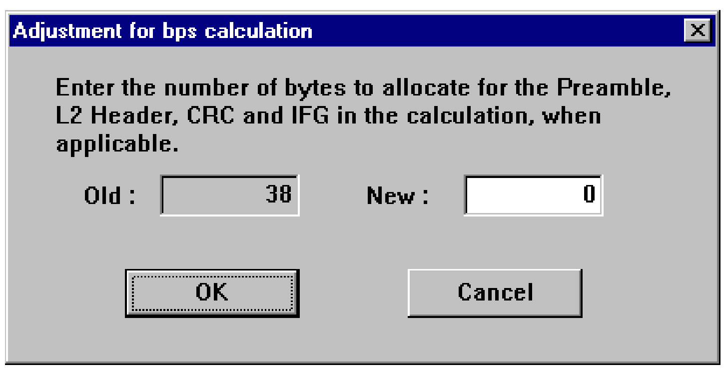

Figure 4-2. Adjustment for bps Calculation Dialog from Detail Tab

To access the *Adjustment for bps calculation* dialog box from the *Detail* tab in test results, click the heading for the **Rx L2 bps** column.

If you want the calculation to be based on the addition of a 4-byte VLAN tag, for example, enter 42.

Note: If you adjust the calculation to 0, the Rx L2 bps values are the same as the Rx L3 bps column values.

Detail Results Navigation Buttons and Line Information

The top of each *Detail* results tab contains navigation buttons to quickly move through results and page number information, as shown in *Figure 4-3*.

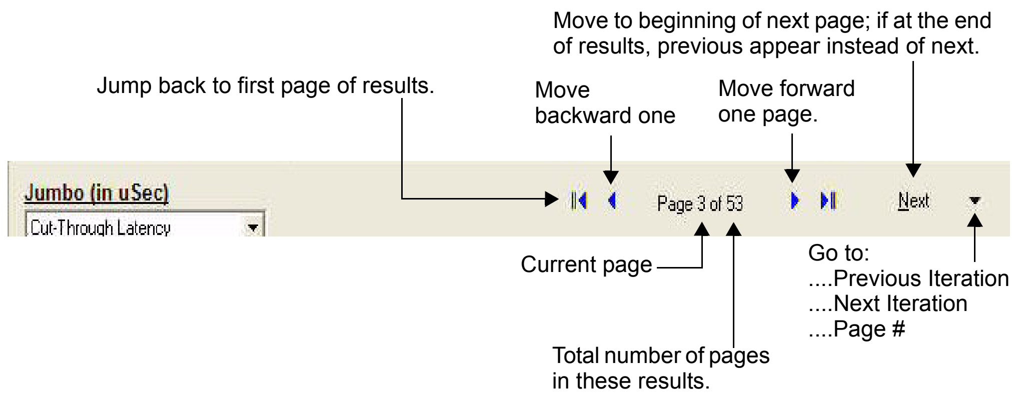

Figure 4-3. Navigation Buttons and Line Count in Detail Results

The number on the right is the maximum number of lines that can show on a page, and it is always 200. The middle number is the total number of rows in the report. The number on the left is the number of the top row of the current view. For example, if you are on page 2 of a 500-row report, the number information would look like: 201-[500]-200.

Table 4-1 shows the buttons on the *Detail* results tabs and their functions.

Table 4-1. Detail Results Navigation Buttons

Button	Function
	Jumps to the first page of the detail results.
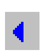	Moves back one page (200 rows) in the detail results.
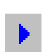	Moves forward one page (200 rows) in the detail results.
	Jumps to the last page of the detail results.

Port Errors Tab

The *Port Errors* test results tab contains the same information from test to test. Port errors vary by card type. For a description of each port error, see *Appendix A, "Port Errors."*

The *Port Errors* tab lists (by port) all Layer 2 port-level errors (such as CRC and alignment) that existed at 2-minute time intervals taken during the test. Errors vary according to the card types (ATM, Ethernet, WAN, or POS) used in the test. *Figure 4-4* shows an example of port errors on the *Port Errors* results tab.

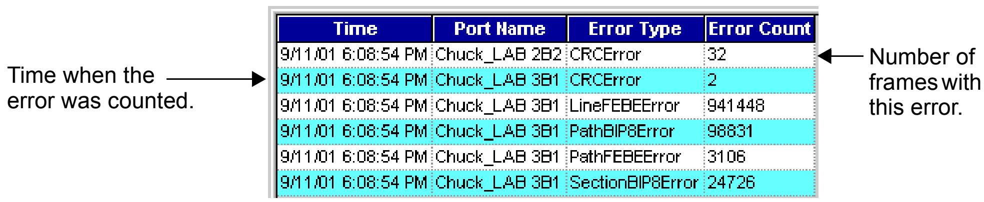

Time	Port Name	Error Type	Error Count
9/11/01 6:08:54 PM	Chuck_LAB 2B2	CRCError	32
9/11/01 6:08:54 PM	Chuck_LAB 3B1	CRCError	2
9/11/01 6:08:54 PM	Chuck_LAB 3B1	LineFEBEError	941448
9/11/01 6:08:54 PM	Chuck_LAB 3B1	PathBIP8Error	98831
9/11/01 6:08:54 PM	Chuck_LAB 3B1	PathFEBEError	3106
9/11/01 6:08:54 PM	Chuck_LAB 3B1	SectionBIP8Error	24726

Figure 4-4. Port Errors Results Tab

Each row in the *Port Errors* tab represents port errors that existed when the count was taken in that test iteration. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

Important: The database holds only one Port Errors table for all errors of all tests such as Throughput, Jumbo, Frame Loss, etc. This means that the results shown on the *Port Errors* tab for all tests will get overwritten by the most recent results of the current test. For example, if a Frame Loss test is run and results are shown on the *Port Errors* tab, and then a Latency test is run with the same setup, the results on the *Port Errors* tab of the Latency test will overwrite the results shown on the *Port Errors* tab for the Frame Loss test.

For a description of each field on the *Port Errors* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Packet Rate Tab

The *Packet Rate* test results tab shows the transmit and receive rate of each port involved in the test for each time sampling interval (in seconds) for each test iteration. You set the frequency at which SmartFlow should take samplings on the *Options* tab. (See *"Packet Rate Options" on page 290*.)

Note: If this tab has no data, you did not select the *Display Tx/Rx packet rate per port* checkbox on the *Options>General* tab.

This tab can display up to 20 ports. If more ports are used in the test, a warning is displayed. Packet counts include only valid packets and no management frames.

Figure 4-5 on page 84 shows an example of packet rates for four ports on the *Packet Rate* results tab. The *Sampling interval (Sec)* field on the *Options>General* tab was set to **5** seconds: SmartFlow takes a packet count sample every 5 seconds, and each time in the *Time Elapsed* column increases by 5 seconds.

For a description of each field on the *Packet Rate* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

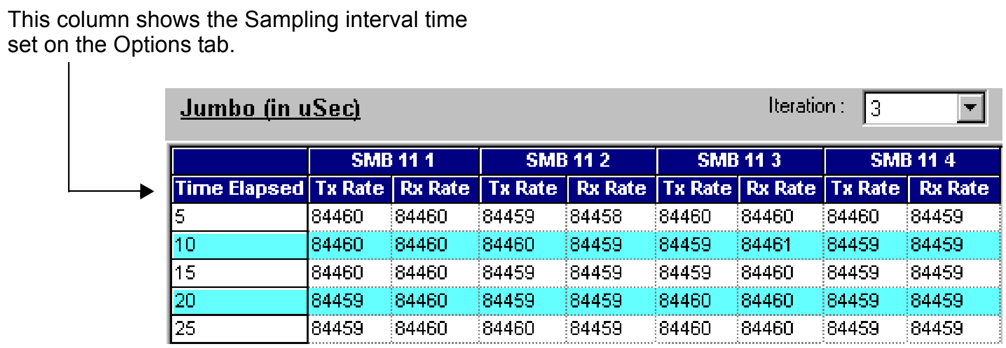

	SMB 11 1		SMB 11 2		SMB 11 3		SMB 11 4	
Time Elapsed	Tx Rate	Rx Rate	Tx Rate	Rx Rate	Tx Rate	Rx Rate	Tx Rate	Rx Rate
5	84460	84460	84459	84458	84460	84460	84460	84459
10	84460	84460	84460	84459	84459	84461	84459	84459
15	84460	84460	84459	84459	84459	84460	84459	84459
20	84459	84460	84459	84459	84460	84460	84459	84459
25	84459	84460	84460	84459	84460	84460	84459	84459

Figure 4-5. Packet Rate Results Tab

Keyboard Shortcuts

You can navigate by using the following key combinations instead of the mouse.

Key Combination	Function
F1	Opens the appropriate online Help for the dialog or window.
F8	Connects to the specified chassis.
Shift + F8	Disconnects from chassis.
Ctrl + F8	Opens the *Setup SmartBits Connections* dialog box for chassis connections.
Ctrl + R	Toggles between results and setup tabs. It switches between the last tabs that were accessed. (This key combination does not include the 0 on the number keypad.)
Ctrl + C	Copies selected test results to the MS Windows clipboard.
Shift + F5	Stops a test in progress.
Alt + F4	Closes SmartFlow.

Basic Actions

This section covers basic actions you need to know to use SmartFlow, such as starting and stopping a test.

Starting a Test and Viewing Results

To start a test and view results:

1. Verify that all selected ports are connected to the DUT.
2. On the shortcut bar, click one of the test launch buttons to start the test. As SmartFlow prepares to run the test, the *Results* window appears. The status bar at the bottom of the *Results* window displays test progress.
3. View the results in chart, summary, or detail format by clicking on the appropriate tab. To change the style of chart, right-click on the chart and select from the pop-up menu. You can also save results to a file.

You can start running the tests in the order in which they appear in the application. Then, depending on results, you can investigate further by rerunning the appropriate tests as required to focus on where the DUT/SUT runs into problems.

You can run any test, however, in any order.

Viewing or Printing Test Results

If you want to print results or view more columns online without scrolling, you can adjust the column widths of the online display. You can adjust the width so that the results print according to your paper size requirements.

Hiding Columns

You can also hide a column by moving the column margin to the left all the way to the previous column margin. To display a hidden column, move the cursor over the previous column heading until it becomes two vertical parallel lines with white in between, and then drag it to the right.

Adjusting Column Size

Place your cursor over side of a column heading and then drag it to make it narrower, wider, or hide it completely, as shown in *Figure 4-6 on page 86*.

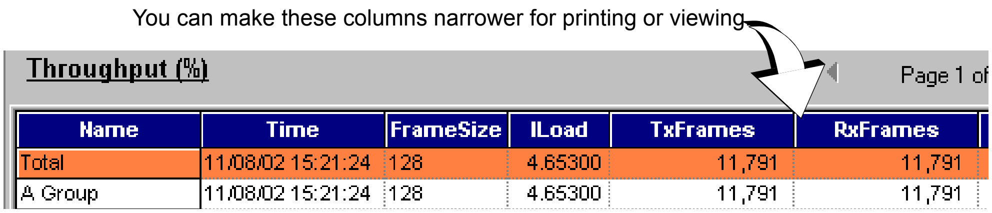

Figure 4-6. Adjust the Column Width Online to Size for Printing or Viewing

Opening a New Configuration File

Whenever SmartFlow is running, a new (unsaved) or existing (previously saved) configuration (`.flo`) file is always open. This file contains the port setup and test parameters, all flows and groups created, and any test results generated during the SmartFlow session. Any changes you make to values exist in the file temporarily until you save it.

To open a new file, select **File>New**.

When you open a new file, SmartFlow prompts you to save any changes to the existing file, then closes the file and opens the new one. The new file defaults to the name *Untitled* until you save it under another name.

Saving Test Configurations and Results

You can save the settings that you specify for ports and tests as well as flows, groups, and any test results by saving them to a configuration (`.flo`) file. Using `.flo` files saves time and is highly recommended. SmartFlow saves the results of the last test or tests run in the configuration file.

Saving a Test Configuration to a File

To save a test configuration to a file, select **File>Save As** and enter the name of the file.

The last saved configuration becomes the default configuration for any future tests.

Important: Make sure to save the configuration *before* you disconnect from the chassis.

Loading a Saved Test Configuration File

SmartFlow allows you to load previously saved `.flo` files. To load a configuration file (saved in a previous session), select **File>Open** and select the file.

If you load the file while you are not connected to a chassis, the saved port configuration is displayed in the *Cards* tab. Once you connect, the current available ports are displayed.

Notes:
- It is strongly recommended that you connect to the chassis in the saved configuration before opening the saved configuration file, especially if you plan to modify it.
- If you open a `.flo` file that was created with a previous version of SmartFlow, once you save it in the current version you can no longer use it in previous versions.

If you save a configuration, reload it, and then try to connect to a chassis that does not match the current file, a warning like the following appears:

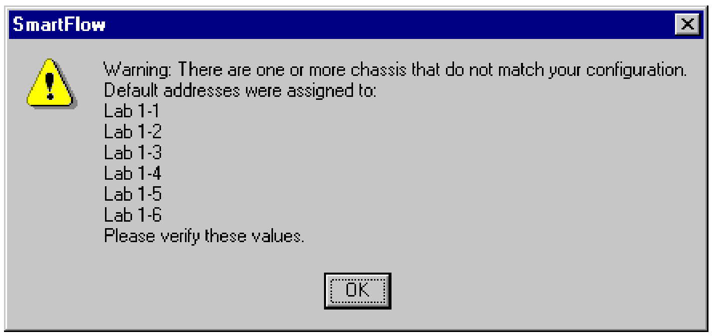

This warning also appears if you configure ports offline, then try to connect to a chassis whose configuration does not match the configuration.

Changing the Tab Display

You can customize the SmartFlow tab display in two ways:

- Show or hide tabs
- Rearrange tabs that are displayed.

Hiding or Showing Tabs

Use the *Customize Tabs* dialog box (*Figure 4-7 on page 88*) to choose which tabs in the main window that you want to show or hide according to your testing needs and preferences.

To access this dialog box from the menu bar, select **Setup>Customize Tabs**.

Tabs for specific protocols or tabs that you may only use infrequently (such as the *Options* tab) may be hidden. For example, if you are testing with IPv6 traffic, you may decide to hide the tabs for ATM, BGP, MPLS LSP, and WAN. Doing this simplifies the window and allows you to work with only the options that pertain to your test.

General tabs that are essential to setting up SmartFlow tests (such as the *SmartFlows* tab or *Test Setup* tab) cannot be hidden.

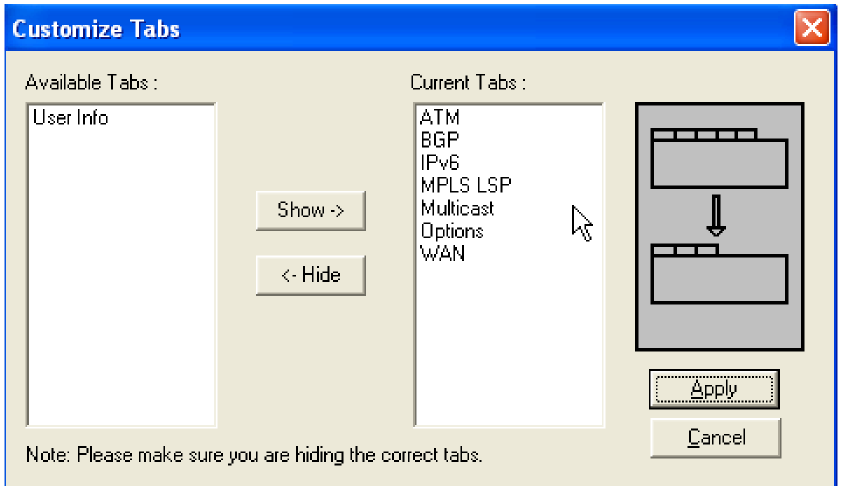

Figure 4-7. Use the Customize Tabs Dialog Box to Show or Hide Tabs

The *Available Tabs* list shows which tabs currently are not displayed. The *Current Tabs* list shows which tabs currently are displayed. Use the *Show->* and *<-Hide* buttons to move tab names from one list to the other.

When you have set the show and hide lists as desired, click **Apply** to enable your selections. Your selections are retained when you reopen the application, even if you did not save the configuration (.flo) file.

To restore the tab display to the default (from the menu bar), select **Setup>Reset Tab Order to Default**.

Rearranging SmartFlow Tabs

You can modify the order of the main SmartFlow tabs to suit your testing needs. (Sub-tabs on the *Test Setup* tab are not modifiable.) This lets you focus on just the technologies that *you* use, by placing those tabs in more prominent positions than tabs for unused or infrequently technologies. For example, if your test doesn't use WAN ports, you do not need to see this tab. You can move the *WAN* tab from the third tab position to after the *MPLS LSP* tab.

Figure 4-8 shows the default order of the tabs.

Figure 4-8. Default Test Setup Tab Order

You can also customize the tabs that are displayed.

To reset the tab order to the default (from the menu bar), select **Setup>Reset Tab Order to Default**.

Figure 4-9 illustrates an example. To move the *WAN* tab to just *before* the *BGP* tab, do this:

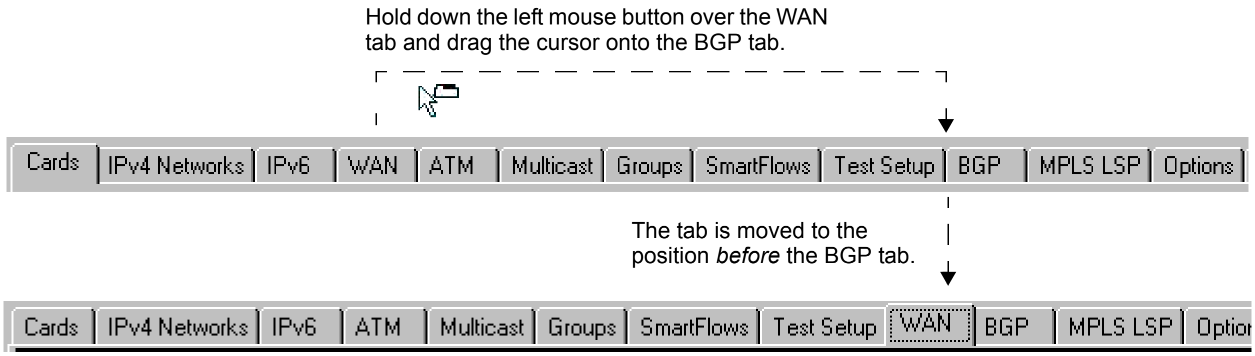

Figure 4-9. Example of Moving a Test Setup Tab

To move a tab to another position:

1. Left-mouse click on the tab that you want to move and hold the mouse button down. A small folder icon appears beside the cursor.
2. Drag the tab *onto* the tab *to the right of* where you want the tab relocated.
3. Release the mouse button.

To move a tab to the very end (right side) of the tab row:

1. Left-mouse click on the tab that you want to move and hold the mouse button down. A small folder icon appears beside the cursor.
2. Drag the tab onto the last tab and release the mouse button. The tab is now in the second from last position.
3. Drag and drop the last tab onto the tab that you actually want to be at the end.

Exporting Test Results to Files

Once you have run a test and results are available, you can export results to an external file(s) for reporting or analysis purposes. You can do any of the following:

- Export results to a `.jpg` (for charts) and `.csv` (comma separated file).
- Export results to HTML.
- Email the contents of any of these results files.

For more information and procedures on each of these export methods, refer to the online Help.

Note: You must first run a test and have results available before you can export results to files. If the export or email options are grayed, make sure that the *Results* view is active by clicking the *Results* button on the shortcut bar on the left side of the window.

You can use any of the options to export or email results based on:

- *The current view.* This exports only the results that are currently displayed on the active results tab. For example, if the *Detail* tab is active when you select the option, only the detail results are exported or emailed.
- *A single test.* This exports all files associated with the test currently being viewed. For example, if you are viewing the Frame Loss test, files for the chart, summary, detail, any stray frames, and port errors associated with that test are exported or emailed.
- *All tests.* This exports all files associated with every test for which results exist in the open.flo file. For example, even if you are viewing the Frame Loss test, all of the files (chart, summary, detail, plus any stray frames and port errors) associated with each test are exported or emailed.

You can also export a test configuration to SAI format. (See *"Exporting a Configuration to SAI" on page 92* and the online Help for more information.)

Setting up Test Information for HTML Export

SmartFlow allows you to export test results to HTML format for portability purposes or to publish your results on the World Wide Web. You can enter summary information for the top of each HTML file. This information includes contact information for your company and the device being tested.

You enter summary user information for a test or all tests at the *User Info* tab, as illustrated in *Figure 4-10*.

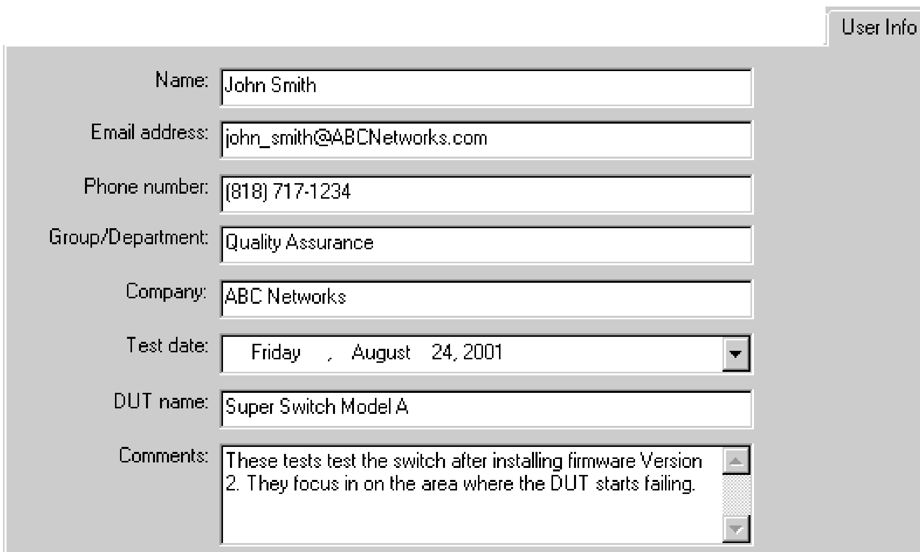

Figure 4-10. User Information Tab

The resulting summary information appears in the `testname.html` and `index.html` files and looks similar to *Figure 4-11*.

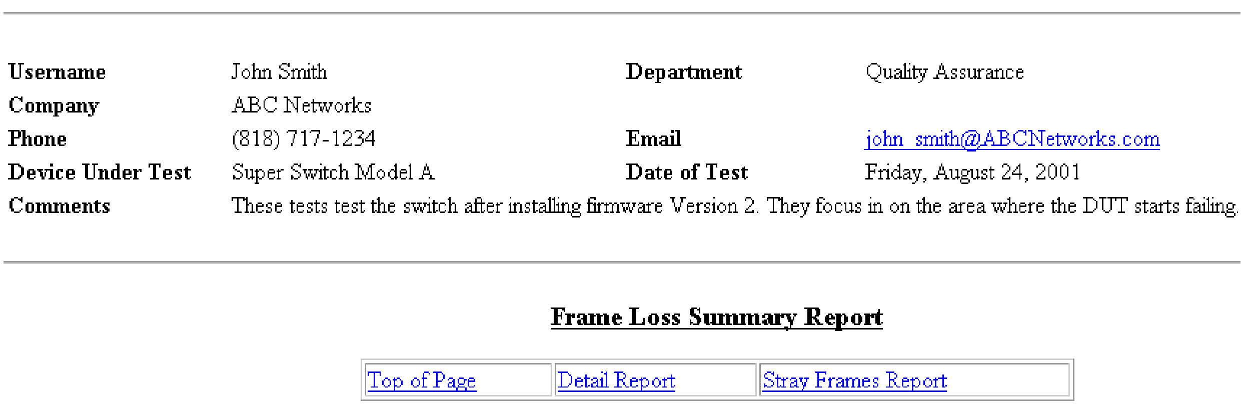

Username	John Smith	**Department**	Quality Assurance
Company	ABC Networks		
Phone	(818) 717-1234	**Email**	john_smith@ABCNetworks.com
Device Under Test	Super Switch Model A	**Date of Test**	Friday, August 24, 2001
Comments	These tests test the switch after installing firmware Version 2. They focus in on the area where the DUT starts failing.		

Frame Loss Summary Report

Top of Page	Detail Report	Stray Frames Report

Figure 4-11. How User Information Appears in Results

Note: Since you can export test results to HTML for all tests as well as individual tests, unless you modify the information on this tab, it appears in the summary/chart file for every test. You may want to check this tab before you run any other tests.

You can utilize these fields to suit your needs, and none of the fields are required.

For a description of each field on the *User Info* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Exporting a Configuration to SAI

You can export the current SmartFlow GUI test configuration to a Script Automation Interface (SAI) script format file. The *Export to SAI* option (accessed from the *File* menu) allows you to transition from running tests using the GUI to an environment for test automation. When you execute your Smartflow tests from scripts, you can run them on UNIX/Linux as well as Windows platforms. Using the SAI can also reduce test execution time.

You can export a test configuration only when SmartFlow is connected to the SmartBits test chassis.

The exported SAI file can be run using either:

- SmartBits Automation SAI (formerly called SmartBits API)
- Standalone SmartFlow SAI.

The SmartFlow GUI automatically determines which type of export file to generate, based on the exported GUI parameters. If all of the exported GUI parameters are supported on both SAI tools, the SAI file created uses the SmartBits Automation SAI format.

Refer to the SmartFlow online Help for information about how to use this feature and the GUI options that are exported. Select **Help** from the menu bar.

Changing the Format of Chart Test Results

Once test results are available, you can view the chart results in either a three-dimensional bar (Manhattan) or ribbon format.

To change the chart type displayed to bar or ribbon:

1 Right-click anywhere on the chart. A popup menu displays.

2 Select **Manhattan** to change to a bar chart or **Ribbon** to change to a flat format.

For a Throughput test, additional options exist for chart and summary results units such as bits per second or frames per second. For more information about these chart and summary options, see *“Chart and Summary Throughput Unit Options” on page 321*.

Stopping a Test in Progress

You can stop a test at any time, even if it is running.

To stop a test while it is running, click the **Stop** button on the toolbar.

Closing SmartFlow

To close SmartFlow, stop any test underway and then do one of the following:

- Click the **Close** button at the top right corner of the screen.
- Select **File>Exit** from the main window menu bar.
- Press the **Alt** + **F4** keys.

Viewing Version Information

To check the version of SmartFlow that you are running or the firmware versions of the cards and chassis to which you are connected, select **Help>About SmartFlow** from the menu bar. The *About SmartFlow* dialog box opens as shown in *Figure 4-12*.

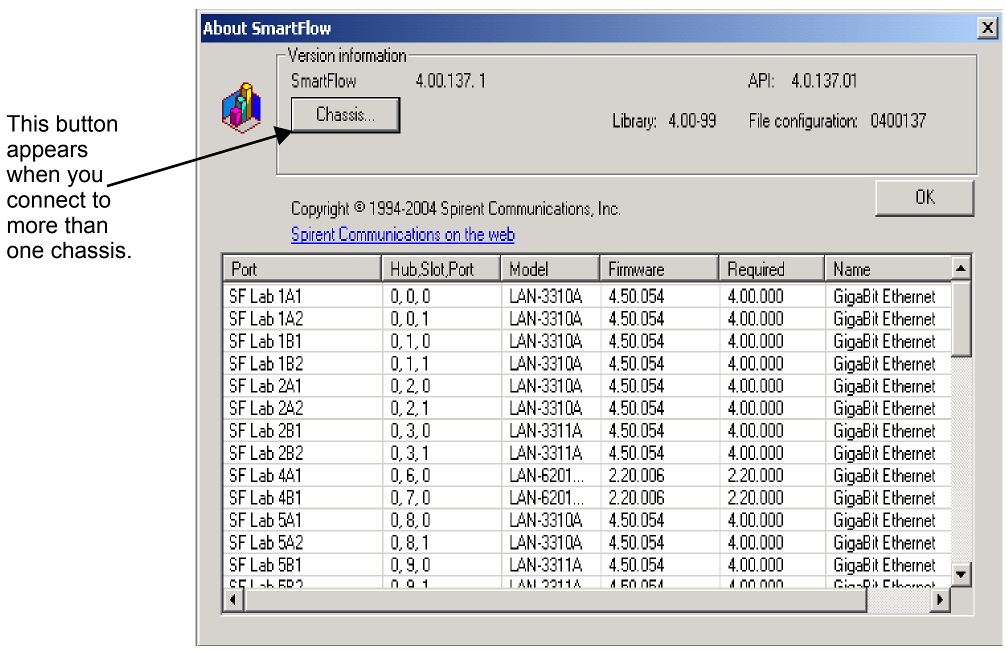

Figure 4-12. About SmartFlow Box Displays Firmware Versions for Chassis and Cards

To check the firmware versions of *multiple chassis* to which you are connected, click the **Chassis** button on the *About SmartFlow* dialog box. The *About Firmware Versions* dialog box opens (*Figure 4-13 on page 94*.)

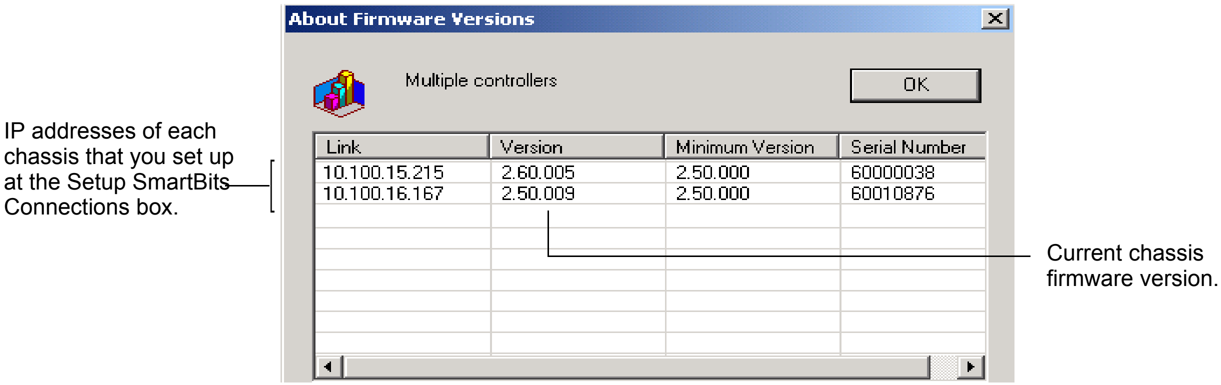

Figure 4-13. About Firmware Versions Box Displays Firmware of Multiple Chassis

Refer to the online Help for descriptions of each field in these dialog boxes.

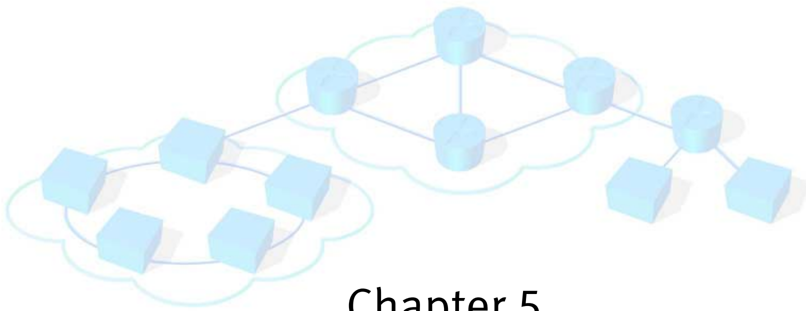

Chapter 5

Set up Test Ports

This chapter provides information about multi-user chassis, SmartBits 600x/6000x port numbering conventions, and how to set up SmartBits test ports.

In this chapter...

Chassis Port Numbering Conventions

SmartFlow displays SmartBits port numbers according to the chassis type. Since SmartFlow allows you to connect to multiple types of chassis (such as a SmartBits 200 and a SmartBits 6000x) simultaneously, the port numbering can vary within the same list of ports.

SmartFlow displays ports in the following places: *IPv4 Networks* tab, *IPv6* tab, *Cards* tab, the *Group Wizard*, *ATM* tab, *WAN* tab, and the *SmartFlows > Traffic* tab. Any port list shows all available ports on the chassis to which you are currently connected.

Notes:
- SmartFlow allows you to modify the port name on the *Cards* tab to a name/number of any length. Keep in mind that long names may be cut off in displays with port names (for example, *Traffic* tab, *WAN* tab). The corresponding column on the *IPv4 Networks* tab is also changed.
- If you change the name of a port, names of flows previously created on that port retain the original name that existed when the flow was created. You can manually change the flow name at the *SmartFlows* tab.

All ports are listed first by the chassis name entered in the *Setup SmartBits Connections* dialog box followed by the port number. See *"Connecting SmartFlow to SmartBits Chassis" on page 68* for information about setting up chassis connections.

Since you can connect to multiple types of chassis at the same time, port numbering varies by chassis type in a port list. *Figure 5-1* illustrates a port list resulting from connections to a SmartBits 2000, a SmartBits 600x, and a SmartBits 6000x. Notice that the numbering scheme for a SmartBits 6000x is different from that of the SmartBits 2000.

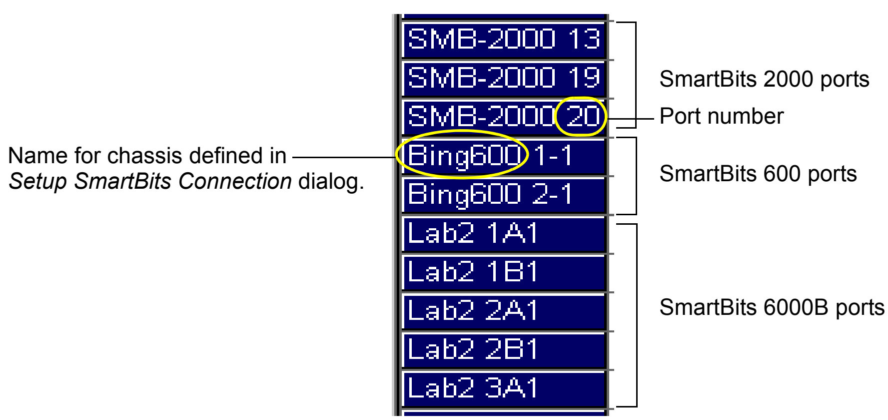

Figure 5-1. Port List with Ports from Multiple Chassis

For more information about port numbering in the SmartBits 600x/6000x, see *"SmartBits 600x/6000x Port Numbering Conventions" on page 97*.

SmartBits 200/2000 Port Numbering Conventions

SmartFlow displays SmartBits 200/2000 port numbers as the chassis name followed by the slot number in the chassis. For example, SmartFlow displays Port 19 in chassis LAB2 (named this when you set up the connection) as:

```
LAB2 19
```

SmartBits 2000 port (slot) numbers start with one and run consecutively from left to right.

SmartBits 200 ports are numbered bottom to top, left to right, as shown here.

SmartBits 200 port numbers

2	4
1	3

SmartBits 600x/6000x Port Numbering Conventions

SmartFlow uses a virtual port numbering system for the SmartBits 600x/6000x and for system status messages (regardless of chassis type) that appear at the bottom of the main window.

SmartBits 6000x Port Numbering

Modules are numbered one through six from top to bottom. Each module is made up of two cards designated A and B. The A card is located in the left side of the chassis and the B card is on the right.

Each port on a module is numbered from left to right. Using this system, an individual port is identified by the module (or slot) number within the chassis, the card, and the port, as illustrated in this excerpt from the *SmartFlows>Traffic* tab.

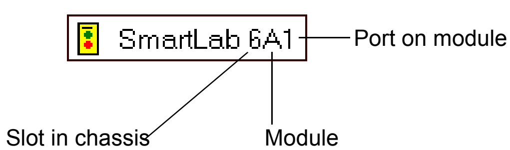

SmartBits 600x Port Numbering

As with the SmartBits 6000x, modules are numbered one to two from top to bottom. However, each module is made up of only one card per slot within the chassis. Each port is numbered from left to right. Using this system, an individual port is identified by the module (or slot) number within the chassis, the card, and the port.

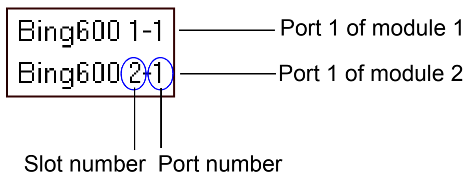

Same-numbered Ports in Multiple Chassis

When the same port number appears on multiple chassis, the chassis name (e.g., LAB 2) appears in the port display and distinguishes between the ports.

However, once flows are created, SmartFlow distinguishes between same-numbered ports on different chassis by prefixing the second and subsequent chassis' ports with a number. (The first chassis has no prefix.)

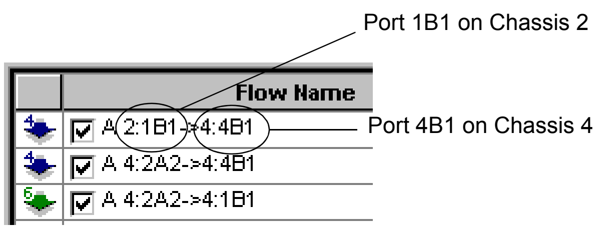

Port Numbering for Ports with Multiple Networks/VLANs/Channels

A port can have multiple networks, and a network can have multiple VLANs. Each subnet or VLAN can be used as a port for a flow. ATM ports can have multiple virtual paths and channels associated with it. WAN ports can have multiple channels and PVCs or PPPs associated with it. SmartFlow lists the ports in various ways, depending on which of these items exist for the port.

If the port has multiple networks, SmartFlow lists the port in one of the ways shown in *Table 5-1 on page 99*.

Table 5-1. Port Display Formats

Port is Listed by...	If...
Network number	• The network is the base network of multiple network(s). • VLAN is disabled for that network.
Network number/VLAN number	VLANs or multiple networks exist for the port and VLAN is enabled.
Port number/VPI/VCI	The port is an ATM port.
Channel/PVC or Channel/PPP name or number	The port is a WAN port.

Ports with multiple networks or VLANs are displayed in the *Group Wizard* as shown in *Figure 5-2*.

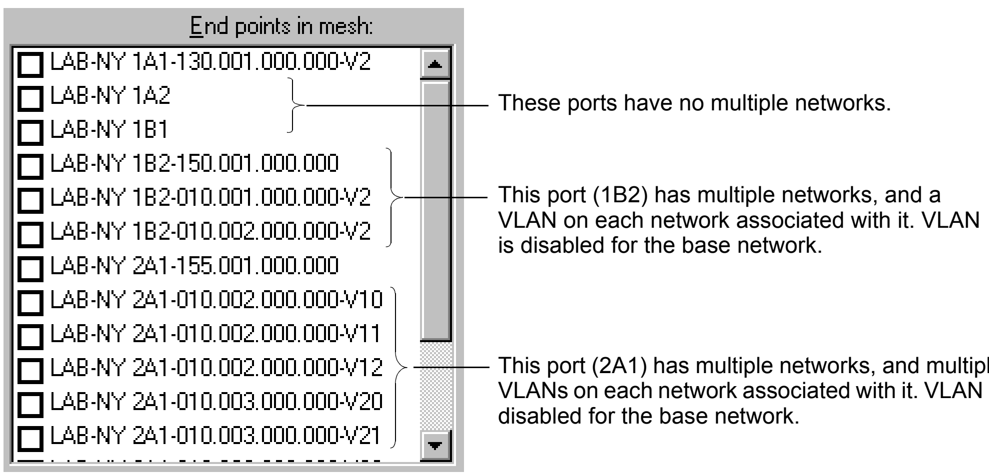

Figure 5-2. Port with Multiple IPv4 Networks

A WAN port with multiple PVCs is displayed in the *Group Wizard* as shown in this example.

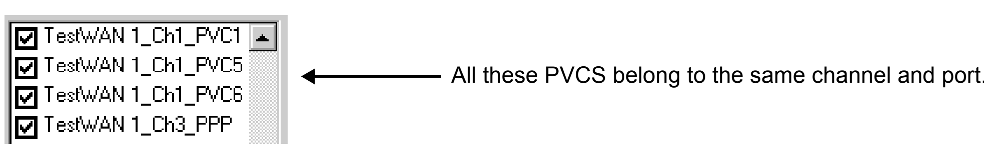

Flow Numbering for Ports with Multiple Networks/VLANs

When a flow is created on a port that has VLANs and/or multiple networks, in the flow list the flow is displayed by the number of the network and then the VLAN ID with which the flow is associated.

The base network is -0, and additional networks are displayed starting with -1. For example, port 4B1-1 indicates the second network (or first subnet) associated with port 4B1, as shown here.

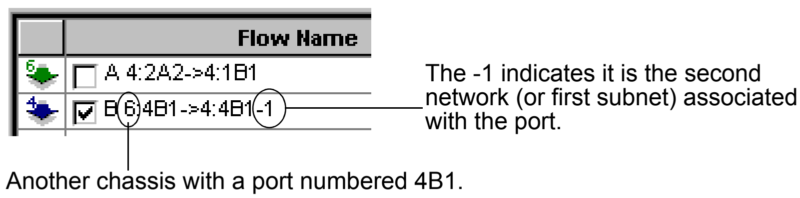

A flow that uses a VLAN on a subnet is displayed by the VLAN ID appended after the subnet index for the appropriate transmitting or receiving port. In the example that follows, each flow was created on the same port.

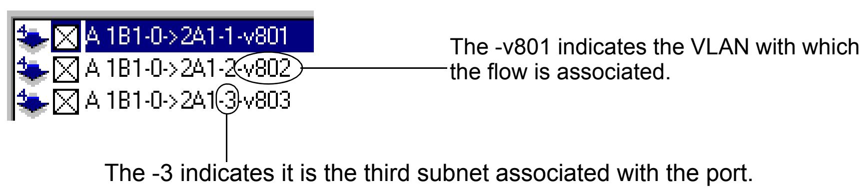

Note: A hyphen and number is also used in flow names when multiple flows are created on ports with no added networks. The number is not based on priority. The hyphen and number are appended after the receiving port in order to uniquely identify each of the multiple flows. For example, the flow C 1A2->1B1-8 is the eighth flow created in a series of flows.

Using Ports in Multi-user Chassis

The information in this section does not apply if you are using a single-user chassis that does not support multiple users.

Multi-user chassis allow multiple users to connect to the same chassis simultaneously, as long as they each use different cards in that chassis. While connected to a multi-user chassis, users can reserve available ports and:

- Run multiple applications
- Run multiple instances of SmartFlow
- Run any other Spirent Communications application.

When you connect SmartFlow to chassis, it displays all cards in the chassis that are supported by SmartFlow. These ports appear on the *Cards* tab that you use to configure and reserve the ports. When you first connect to a chassis with multi-user capabilities enabled (SmartBits 2000 or SmartBits 600x/6000x), only the ports that are not reserved by anyone else and are available appear in the list.

Note: Cards display only if the correct firmware is loaded. Go to the Spirent Communications web site at http://www.spirentcom.com to download the latest firmware.

Displaying Available Ports in a Multi-user Chassis

When you connect to a multi-user chassis, all cards in the chassis that are supported by SmartFlow are displayed in the *Cards* tab, the G*roup Wizard*, and the *SmartFlows>Traffic* tab. If another user is using the port, it displays grayed with a status of *In Use*. with the *Reserve* field already checked.

The following shows possible states of a port in a multi-user chassis on the *Cards* tab:

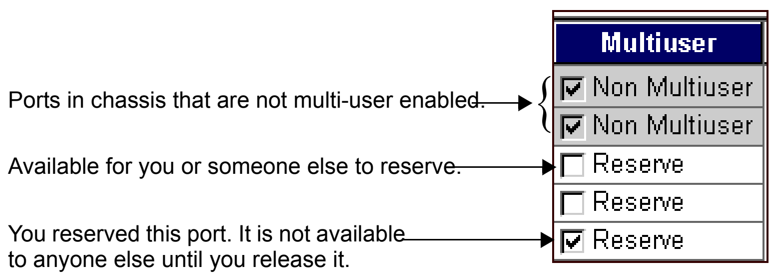

If it has been a while since you connected to a multi-user chassis but you have not yet reserved any ports, it is recommended that you refresh the display since it is possible that another user has reserved ports.

To refresh the available ports display:

1. Click the heading of the **Multiuser** column to highlight the entire column.
2. Right-click anywhere on the column while it is highlighted. A popup menu appears.
3. Select **Refresh**. If anyone else reserved or released a card since you connected, its status changes to *Reserve* unchecked or *In Use,* respectively.

Reserving Ports in a Multi-user Chassis

To ensure that the port you have selected does not get reserved by another user prior to running your test, reserve the port by selecting the checkbox in the **Multiuser** column on the *Cards* tab. While it is reserved it does not appear as available to anyone else either using another copy of SmartFlow or any other Spirent Communications application.

Note: You can run a test without manually reserving the ports. SmartFlow automatically reserves the test ports while the test is running and automatically releases them when the test finishes. However, by not reserving the port prior to the test, the test may fail if someone else reserved the port before you tried to start it.

Reserving Ports on Cards with Multiple Ports

When you reserve a port on a module containing multiple ports (such as the LAN-3101A/B), you actually reserve *all* of the ports on the card (PCB) containing the port that you reserved even if you are not using all of them. *Figure 5-3* illustrates this situation.

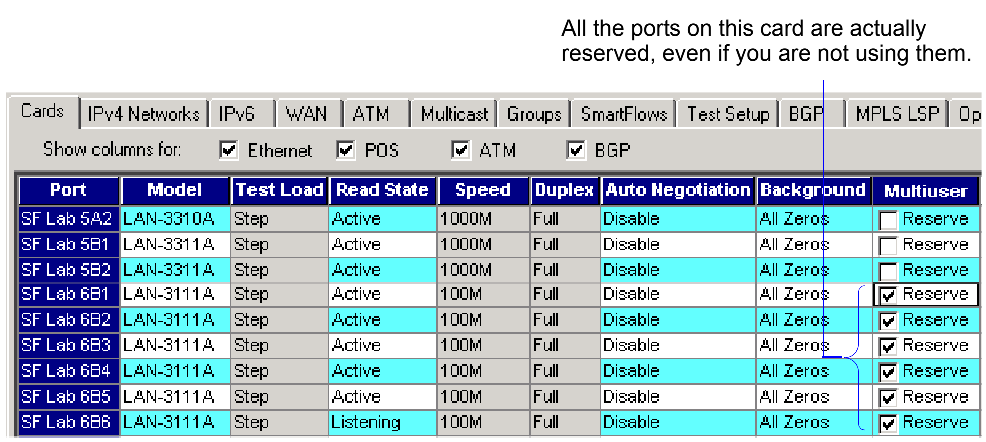

Figure 5-3. Cards with Multiple Ports on Cards Tab

If you reserved only *one* port on a module containing multiple ports, when you release it the remaining ports on the card are also released. If you reserved more than one port, you must release *each* of the reserved ports to release the rest of the ports on the card.

What Happens if Someone Reserves the Port Before I Do?

The ports listed on the *Cards* tab, *SmartFlows>Traffic* tab, and *Group Wizard* reflect ports that are available at the time you connected to the chassis. However, only the *Cards* tab reflects the port's current availability. Since you connected, if someone else reserves a port before you do, the status changes to *In Use*.

When you try to run a test with *un*reserved ports without first checking their availability, if someone else reserved one (or more) of the ports prior to the start of the test, you receive a message for each unavailable port such as:

```
Unable to Reserve LAB2 1B1
```

You can either select another port, omit that port from the test, or wait until the port is available.

Unreserving a Port in a Multi-user Chassis

If you reserve and configure a port and then decide not to use it in a test, it is recommended that you unreserve the port to make it available for other users.

To unreserve a port:

1. Select the **Cards** tab.
2. Clear the checkbox in the **Multiuser** field for the port in order to unreserve it.

Configuring Test Ports

Before you run a test or set up flows, configure the ports involved in the test at the *Cards* tab on the main window. This is an important step, because SmartFlow derives the flow's Layer 2 address from this information. Use the *Cards* tab to configure the card's physical and link level (MAC address) port parameters. To specify additional networks, use the *IPv4 Networks* tab.

You can send traffic between ports that are both configured for IPv4 or IPv6. You can also send traffic between a port configured for IPv4 and a port configured for IPv6, and vice versa. See *"Defining IPv6 Information for a Port" on page 112* for more information.

Important: For ports in a multi-user chassis, until you connect and reserve the port in the *Multiuser* column on the *Cards* tab (*Figure 5-4*), another user can reserve the card.

Once you connect to a SmartBits chassis, the *Cards* tab displays all ports in each chassis to which you are connected. Scroll to the right to display the remaining fields on the tab.

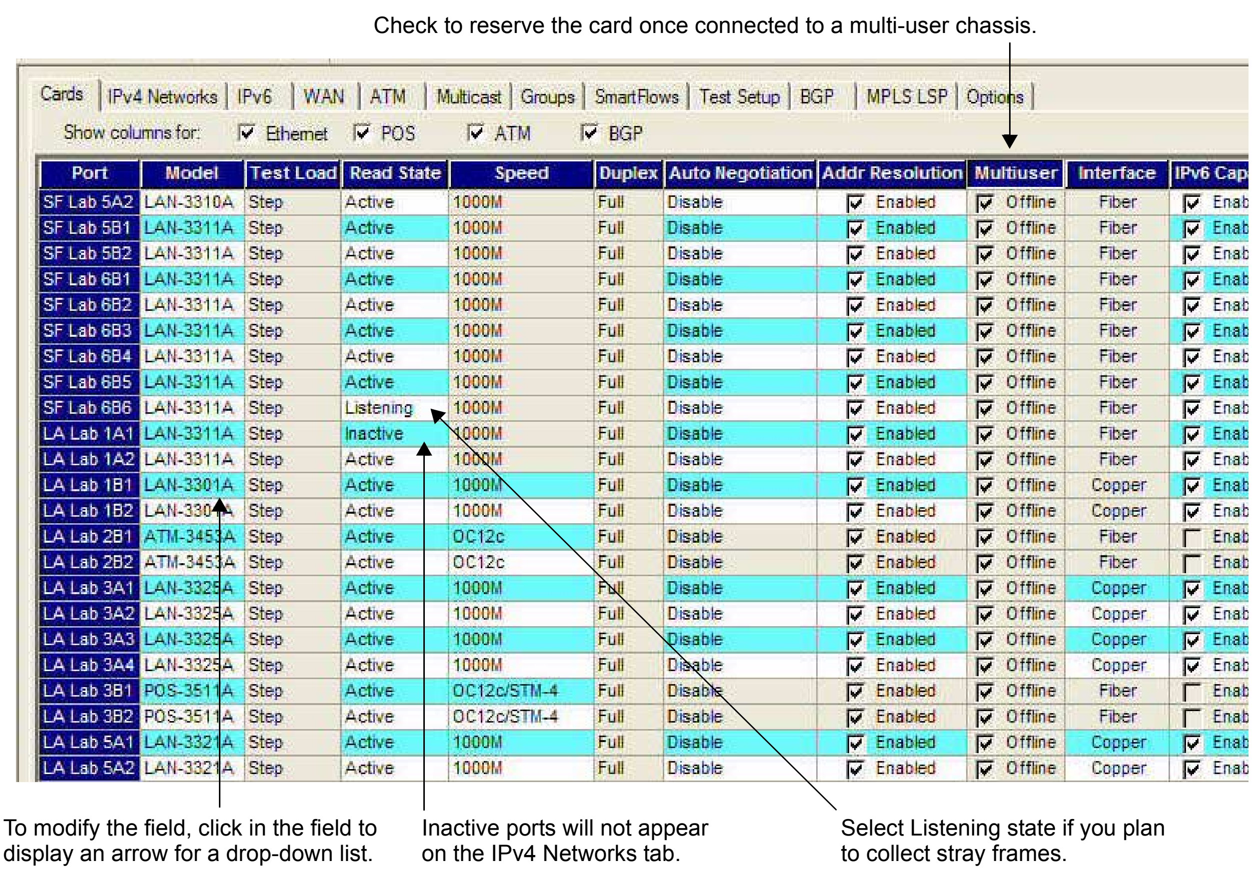

Port	Model	Test Load	Read State	Speed	Duplex	Auto Negotiation	Addr Resolution	Multiuser	Interface	IPv6 Cap
SF Lab 5A2	LAN-3310A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 5B1	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 5B2	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 6B1	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 6B2	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 6B3	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 6B4	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 6B5	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
SF Lab 6B6	LAN-3311A	Step	Listening	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
LA Lab 1A1	LAN-3311A	Step	Inactive	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
LA Lab 1A2	LAN-3311A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Fiber	☑ Enab
LA Lab 1B1	LAN-3301A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 1B2	LAN-3301A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 2B1	ATM-3453A	Step	Active	OC12c	Full	Disable	☑ Enabled	☑ Offline	Fiber	☐ Enab
LA Lab 2B2	ATM-3453A	Step	Active	OC12c	Full	Disable	☑ Enabled	☑ Offline	Fiber	☐ Enab
LA Lab 3A1	LAN-3325A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 3A2	LAN-3325A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 3A3	LAN-3325A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 3A4	LAN-3325A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 3B1	POS-3511A	Step	Active	OC12c/STM-4	Full	Disable	☑ Enabled	☑ Offline	Fiber	☐ Enab
LA Lab 3B2	POS-3511A	Step	Active	OC12c/STM-4	Full	Disable	☑ Enabled	☑ Offline	Fiber	☐ Enab
LA Lab 5A1	LAN-3321A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab
LA Lab 5A2	LAN-3321A	Step	Active	1000M	Full	Disable	☑ Enabled	☑ Offline	Copper	☑ Enab

Figure 5-4. Cards Tab

The alternating colors on the grid simply aid in reading information and distinguishing between ports, especially when you have many networks in your configuration.

Notes:

- For POS cards, make sure that the POS and PPP setup in SmartFlow matches the DUT's setup for POS and PPP.
- Make sure that you connect to chassis before you set up the ports, networks, and flows to ensure that flows and ports belong to the correct networks.
- To set up WAN cards, also use the *WAN* tab.
- To set up ATM ports, also use the *ATM* tab

For a description of each field on the *Cards* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Customizing the Columns Displayed

For easier viewing, you can specify the columns that you want displayed on the *Cards* tab according to technology. By selecting the appropriate checkboxes in the *Show columns for* field at the top of the tab, you can display only the columns for the technology that you are testing. For example, if you do not plan to test ATM, clear the ATM checkbox. You can select one or all of these:

- Ethernet-specific fields
- POS-specific fields
- ATM-specific fields
- BGP-specific fields.

Columns associated with a technology are displayed to the right of existing columns in the order in which the technology options appear on the tab. The order in which you select the checkboxes does not affect the order in which the columns are displayed. For example, ATM fields always appear before (to the left of) POS fields.

Mismatched Port Configurations

Mismatched port configurations can occur if you configure the ports offline. If you try to connect to a chassis and the ports configured offline do not match those in the chassis, a warning box appears alerting you to this situation. See *"Loading a Saved Test Configuration File" on page 86* for more information.

It is recommended that you connect to chassis *before* you set up the ports, networks, and flows. This ensures that flows and ports belong to the correct networks. If you plan to modify a previously saved configuration file, connect to chassis before you load the file.

A warning appears if you save a configuration, reload it, and then try to connect to a chassis that does not match the current file.

Defining IPv4 Network Information for a Port

To make traffic more realistic, you can modify and add networks, gateways, IP addresses, and VLAN IDs for ports by using the *IPv4 Networks* tab. Use the *IPv4 Networks* tab to define all Layer 3 information for the port. To specify physical and link level (MAC address) parameters, use the *Cards* tab.

You can define networks before or after setting up the ports at the *Cards* tab.

Figure 5-5 shows the *IPv4 Networks* tab with multiple networks and VLANs defined for several ports.

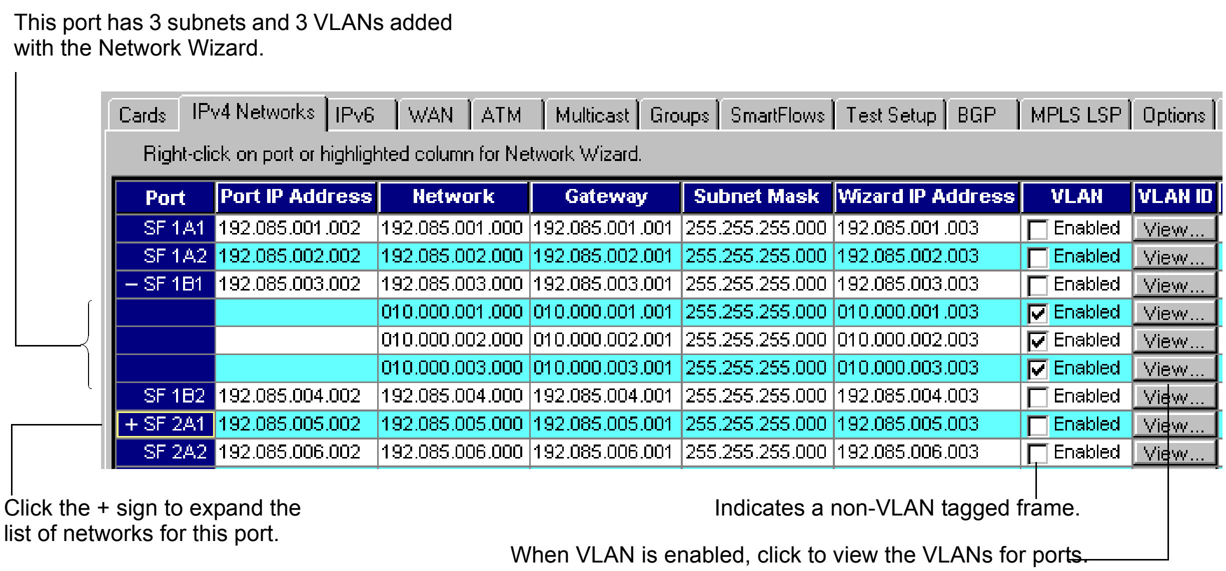

Port	Port IP Address	Network	Gateway	Subnet Mask	Wizard IP Address	VLAN	VLAN ID
SF 1A1	192.085.001.002	192.085.001.000	192.085.001.001	255.255.255.000	192.085.001.003	Enabled	View...
SF 1A2	192.085.002.002	192.085.002.000	192.085.002.001	255.255.255.000	192.085.002.003	Enabled	View...
− SF 1B1	192.085.003.002	192.085.003.000	192.085.003.001	255.255.255.000	192.085.003.003	Enabled	View...
		010.000.001.000	010.000.001.001	255.255.255.000	010.000.001.003	☑ Enabled	View...
		010.000.002.000	010.000.002.001	255.255.255.000	010.000.002.003	☑ Enabled	View...
		010.000.003.000	010.000.003.001	255.255.255.000	010.000.003.003	☑ Enabled	View...
SF 1B2	192.085.004.002	192.085.004.000	192.085.004.001	255.255.255.000	192.085.004.003	Enabled	View...
+ SF 2A1	192.085.005.002	192.085.005.000	192.085.005.001	255.255.255.000	192.085.005.003	Enabled	View...
SF 2A2	192.085.006.002	192.085.006.000	192.085.006.001	255.255.255.000	192.085.006.003	Enabled	View...

Figure 5-5. IPv4 Networks Tab

The alternating colors on the grid are an aid to reading information and distinguishing between ports, especially when you have many networks in your configuration.

To modify default values:
Click in the field to highlight it and enter the new value. You can modify multiple cells in a column the same way you do on the *Cards* tab, by using the right-click menu. If you have defined multiple subnets, however, do not use the *Copy Down* function to make the subnets identical on each port (e.g., same values in *Port IP Address*, *Network,* and *Gateway* fields). This will cause incorrect test results.

Set field values working from left to right. For tests that use ports *all within the same subnet*, first configure the *Subnet Mask*, then change the *Gateway* and *Network* columns. This ensures that the fields on the tab update correctly.

Be sure to connect to the chassis before you set up ports, networks, and flows. This ensures that flows and ports belong to the correct networks.

Notes:

- The *IPv4 Networks* tab does not apply to WAN or ATM cards. These cards do not appear on the tab even if they are in your chassis. To set up WAN cards, use the *WAN* tab. To set up ATM cards, use the *ATM* tab.
- For multicast traffic, the *IPv4 Networks* tab fields apply only to the transmitter of the group. To set up VLANs for receivers of the group, use the *Multicast* tab. See *"Defining Multicast Groups and VLANs" on page 475*.

For a description of each field on the *IPv4 Networks* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar. The Help also provides usage information such as how to modify multiple cells on the tab.

Defining VLAN IDs and Subnet IDs

Use the *VLAN* and *VLAN ID* columns to add VLANs to the networks. To enable VLAN for a network, select the **Enabled** checkbox. This enables the *View* button in the *VLAN ID* column. The *View* button opens the *VLAN ID Editor* dialog box, which is used to set VLAN IDs.

Set up multiple VLANs quickly, on the same or other ports, by using a right-click dynamic menu on the *IPv4 Networks* tab and *IPv6* tab. In the *VLAN* column, select multiple cells by using the mouse and Shift or Control keys. Then right-click to open the dynamic menu (*Figure 5-6*). For the selected networks, it enables you to:

- Enable or disable VLAN.
- Copy down VLAN IDs.
- Increment or decrement the VLAN IDs.
- Remove VLAN IDs.
- Copy VLAN Subnet IDs.

For complete information about how to use these options, refer to the SmartFlow online Help.

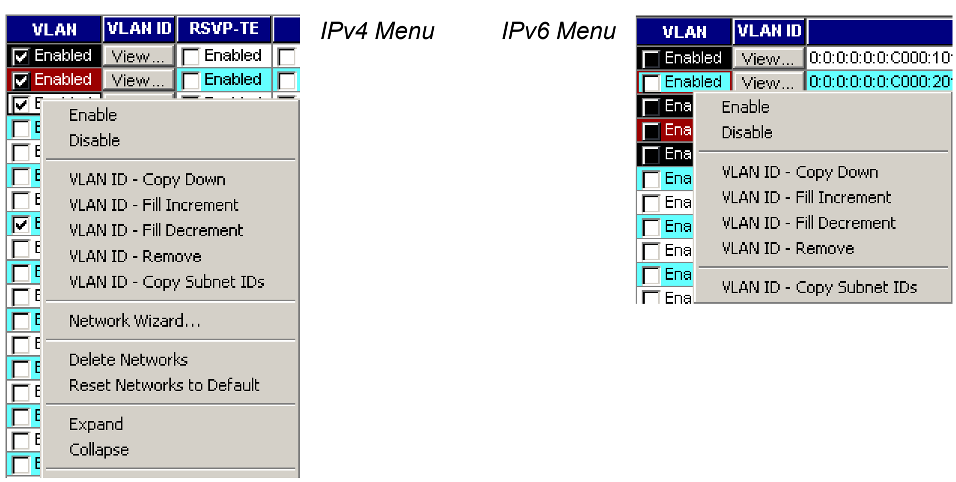

Figure 5-6. Dynamic Menu for VLAN Definitions

Associating Multiple Networks and VLANS with Ports

You can use the Network Wizard for IPv4 and the IPv6 Network Wizard to add multiple networks (gateways or subnets) and VLANs to a port. Doing this makes the port behave more like a part of a network cloud instead of an edge device. You can add multiple networks, with or without VLAN IDs, to one or more selected ports. If you wish to add only VLANs to an existing network, but not add networks, use the *VLAN ID Editor* dialog box. (See *"Defining VLAN IDs and Subnet IDs" on page 107*.)

To open the appropriate *Network Wizard*, right-click over a port number in the *IPv4 Networks* or *IPv6* tab (*Figure 5-7*).

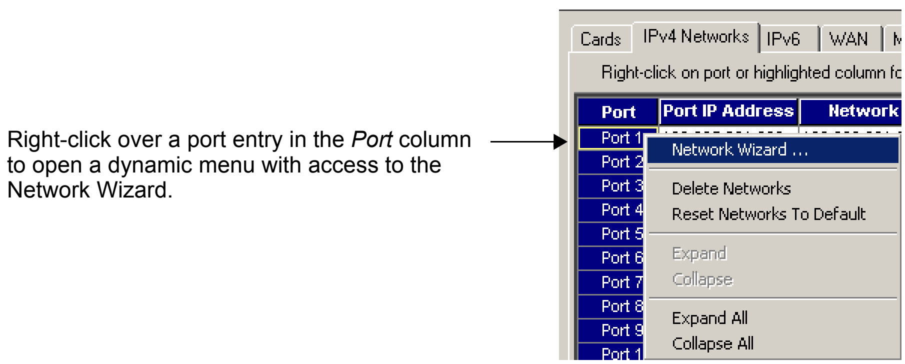

Figure 5-7. Opening the Network Wizard

The Network Wizard associates multiple networks and VLANs with one or more ports. It simultaneously adds and increments networks to the ports that you highlight on the *IPv4 Networks* or *IPv6* tab.

- For IPv4, it bases new subnet IP addresses on the subnet mask and gateway in the *Network Wizard* dialog box. The Network Wizard automatically increments the network, gateway, and IP address for the number of networks that you want to add.
- For IPv6, it bases the new subnets on the prefix length and the interface IP address in the *IPv6 Network Wizard* dialog box.

Figure 5-8 on page 109 shows the Network Wizard for IPv4.

By default, each network has one VLAN ID if you enable VLANs on the *IPv4 Networks* tab or *IPv6* tab. The Network Wizard increments the VLAN ID by one for each VLAN added.

If you associate multiple networks with a port, you must configure the DUT with the proper routing information before running the test, since this information is not dynamically updated.

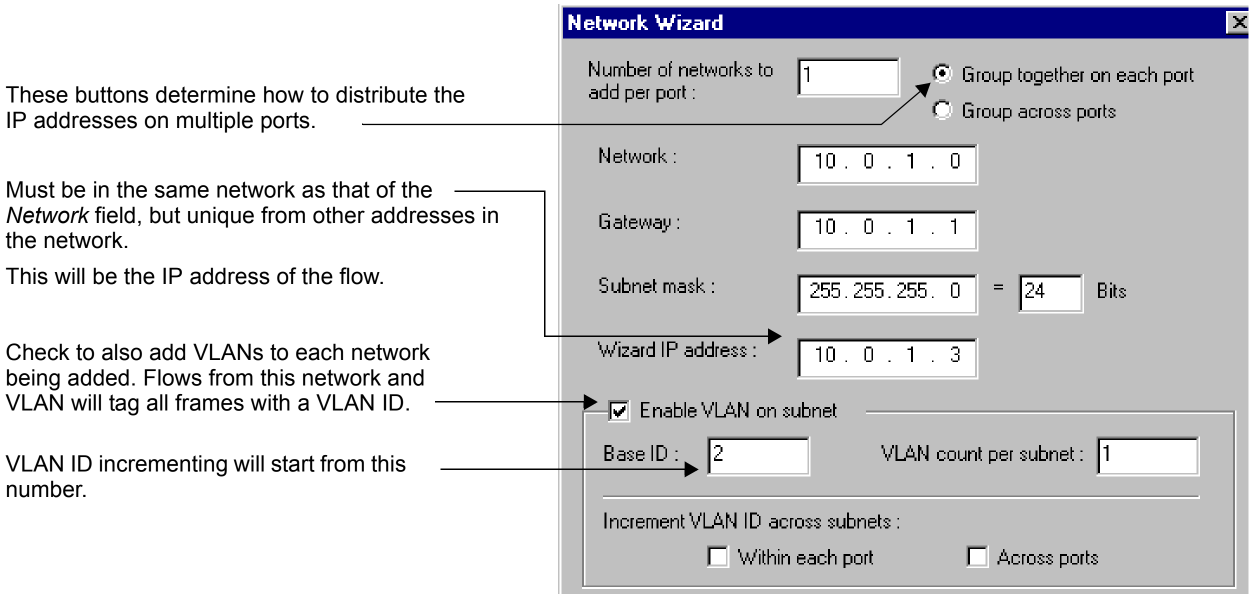

Figure 5-8. Associate Multiple Networks and VLANs Using the IPv4 Network Wizard

Notes:

- You can set up networks with and without VLAN tags on the same port, but you must assign them to different network numbers on that port.
- If you set up multiple networks on the same port, each network must have a unique address.
- If you set up multiple VLANs on networks, each VLAN ID must be unique per network.
- Only Ethernet ports support multiple VLANs.
- When you disable VLAN, existing VLANS are not deleted. They are disabled and stored.

Note: It is highly recommended that you use the default IP address for the *Wizard IP address* field.

For additional information about the Network Wizard, see the online Help. It provides field descriptions and explains how to add or delete networks, add or modify VLANs on networks, and keep the same MAC address with multiple networks on a port.

As an example, *Figure 5-9* shows the settings to add two networks to each highlighted port. This example includes adding two networks to four ports using the *Group together on each port* option.

Based on this example, *Figure 5-10 on page 111* shows the networks (and VLANs) that would result.

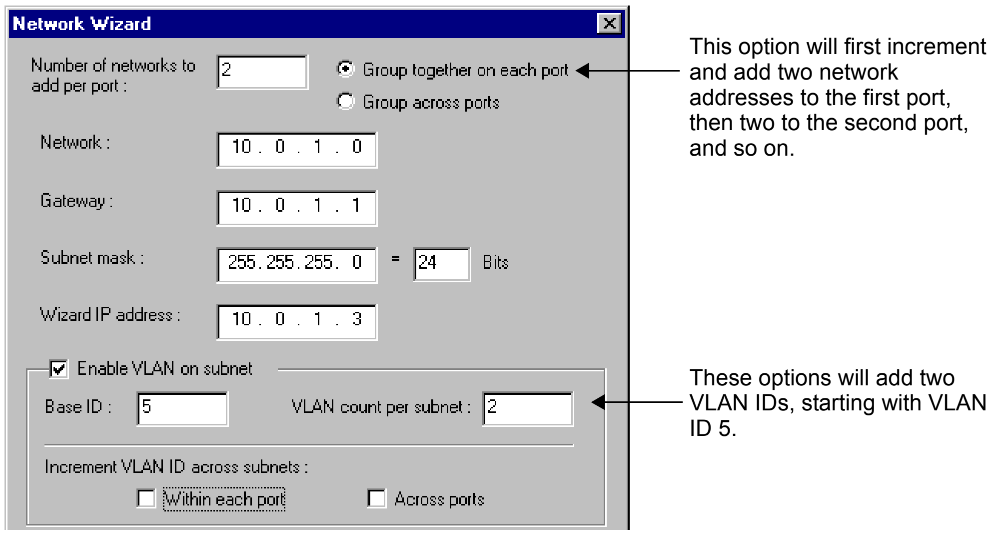

Figure 5-9. Setup to Add Two Networks to Each Highlighted Port

When you set up flows, the source or destination IP address of the flow must be part of one of the networks associated with the respective port on the *IPv4 Networks* tab.

Notice how two networks and VLANs were incremented for one port, then for the next.

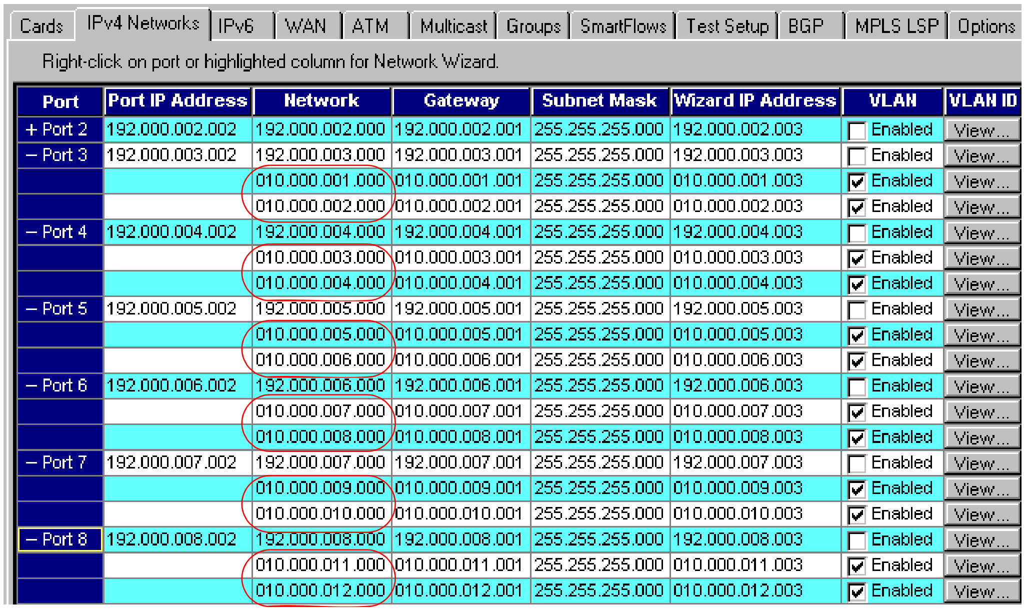

Port	Port IP Address	Network	Gateway	Subnet Mask	Wizard IP Address	VLAN	VLAN ID
+ Port 2	192.000.002.002	192.000.002.000	192.000.002.001	255.255.255.000	192.000.002.003	☐ Enabled	View...
– Port 3	192.000.003.002	192.000.003.000	192.000.003.001	255.255.255.000	192.000.003.003	☐ Enabled	View...
		010.000.001.000	010.000.001.001	255.255.255.000	010.000.001.003	☑ Enabled	View...
		010.000.002.000	010.000.002.001	255.255.255.000	010.000.002.003	☑ Enabled	View...
– Port 4	192.000.004.002	192.000.004.000	192.000.004.001	255.255.255.000	192.000.004.003	☐ Enabled	View...
		010.000.003.000	010.000.003.001	255.255.255.000	010.000.003.003	☑ Enabled	View...
		010.000.004.000	010.000.004.001	255.255.255.000	010.000.004.003	☑ Enabled	View...
– Port 5	192.000.005.002	192.000.005.000	192.000.005.001	255.255.255.000	192.000.005.003	☐ Enabled	View...
		010.000.005.000	010.000.005.001	255.255.255.000	010.000.005.003	☑ Enabled	View...
		010.000.006.000	010.000.006.001	255.255.255.000	010.000.006.003	☑ Enabled	View...
– Port 6	192.000.006.002	192.000.006.000	192.000.006.001	255.255.255.000	192.000.006.003	☐ Enabled	View...
		010.000.007.000	010.000.007.001	255.255.255.000	010.000.007.003	☑ Enabled	View...
		010.000.008.000	010.000.008.001	255.255.255.000	010.000.008.003	☑ Enabled	View...
– Port 7	192.000.007.002	192.000.007.000	192.000.007.001	255.255.255.000	192.000.007.003	☐ Enabled	View...
		010.000.009.000	010.000.009.001	255.255.255.000	010.000.009.003	☑ Enabled	View...
		010.000.010.000	010.000.010.001	255.255.255.000	010.000.010.003	☑ Enabled	View...
– Port 8	192.000.008.002	192.000.008.000	192.000.008.001	255.255.255.000	192.000.008.003	☐ Enabled	View...
		010.000.011.000	010.000.011.001	255.255.255.000	010.000.011.003	☑ Enabled	View...
		010.000.012.000	010.000.012.001	255.255.255.000	010.000.012.003	☑ Enabled	View...

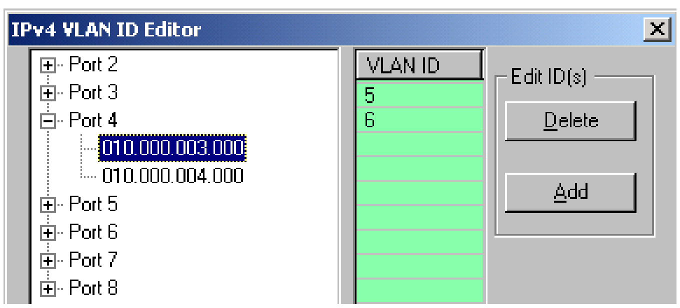

In this case, these VLANs will be added to each port that was highlighted.

Figure 5-10. Example of Resulting Networks and VLANs with Group together on each port Option

Defining IPv6 Information for a Port

For a port to send or receive IPv6 traffic, you must enable IPv6 for the port on the *Cards* tab and define the IPv6 addresses for the port. You define IPv6 address information for ports at the *IPv6* tab (*Figure 5-11*).

The tab displays only the IPv6-capable ports. Modules that support IPv6 include the following: LAN-31xx, LAN-33xx, XLW-372x, and XFP-373x.

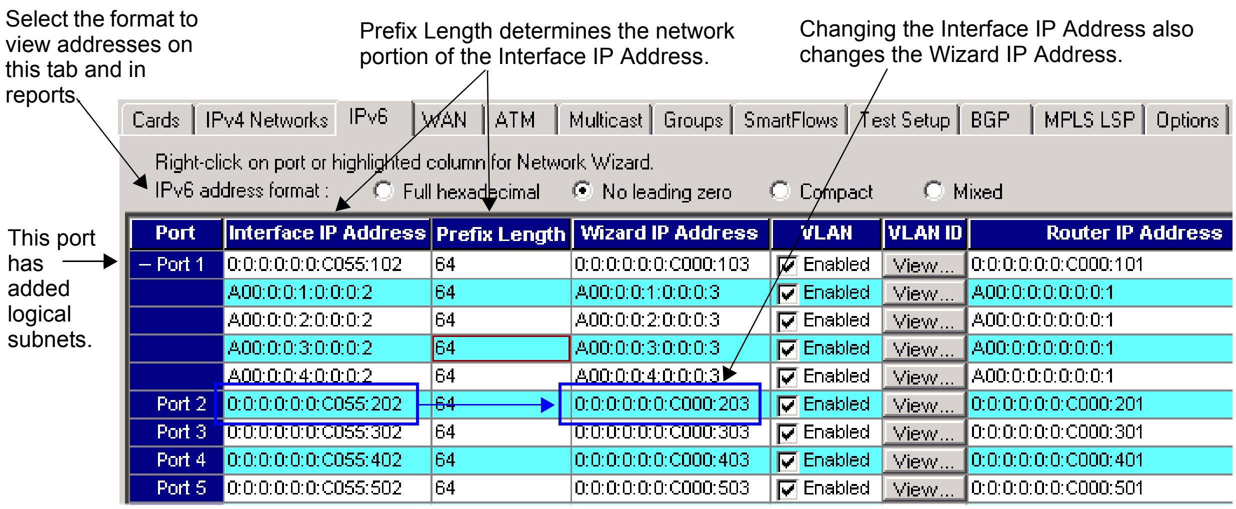

Figure 5-11. IPv6 Tab

To make traffic more realistic, you can modify and add networks and VLAN IDs to a port. To enable VLAN for a network, select the **Enabled** checkbox. This enables the *View* button in the *VLAN ID* column. To add VLANs, click the **View** button under the *VLAN ID* column. This opens the *IPv6 VLAN ID Editor* dialog box. Use it to view, add, or delete VLANs from any network on ports displayed on the *IPv6* tab.

Use the *IPv6* tab to define all Layer 3 information for the port. To specify physical and link level (MAC address) parameters, use the *Cards* tab. You can define networks before or after setting up the ports at the *Cards* tab.

SmartFlow automatically converts default IPv4 addresses on the *IPv4 Networks* tab to default IPv6 addresses on the *IPv6* tab. It does this by prepending 96 zero bits to the IPv4 address. However, if you modify the default IPv4 address for a port, the corresponding IPv6 address will *not* change.

You can modify multiple addresses in a column on the *IPv6* tab by incrementing, decrementing, randomly filling or copying from the first highlighted address down to the remainder of the highlighted cells.

For more information, see *"Associating Multiple Networks and VLANS with Ports" on page 108* and *"Defining VLAN IDs and Subnet IDs" on page 107*.

IPv6 Address Formats in SmartFlow

The addresses on the *IPv6* tab correspond to the IPv4 equivalent addresses on the *IPv4 Networks* tab. By default, SmartFlow converts the IPv4 address of the port to an IPv6 address by prepending leading zero bits to a hexadecimal equivalent of the IPv4 address. For example, a *Wizard* address in IPv4 of 192.085.001.003 becomes in IPv6 0000:0000:0000:0000:0000:0000:C055:0103.

At the top of the *IPv6* tab, you can select the format in which you want to enter/view IPv6 addresses.

Note: The format that you select for IPv6 addresses is also displayed in this format in other setup tabs (such as the *IP* tab), as well as in test results.

Table 5-2 shows a Wizard IP address from the *IPv4* tab in each of the IPv6 address formats.

Table 5-2. IPv6 Address Format Comparison

Format	Address
IPv4	192.085.001.003
IPv6 *Full hexadecimal*	0000:0000:0000:0000:0000:0000:C055:0103
IPv6 *No leading zero*	0:0:0:0:0:0:C055:103
IPv6 *Compact*	::C055:103
IPv6 *Mixed*	::192.085.001.003

For information about how to set up IPv6 tests, see *"IPv6 Test Setup" on page 298*.

For a description of each field on the *IPv6* tab, refer to the SmartFlow online Help. Click **F1** over the tab, or select **Help** from the menu bar. The online Help also describes how to modify multiple addresses and other cells on the *IPv6* tab.

IPv6 Extension Headers

IPv4 allows various option fields to be inserted into a frame. These options make possible special handling for certain packets, such as requesting a specific route (source routing). When these option fields are present, extra processing by nodes is required. As a result, performance is usually reduced.

In IPv6, the option fields are removed. In their place, extension headers can be inserted into frames for the same purpose: to specify optional internet-layer information. The extension headers are separate headers that may be inserted into a packet between the IPv6 header and the upper-layer header (such as TCP). In addition, IPv6 extension headers for IP, TCP, and UDP flows can be edited.

Note: IPv6 extension headers are only supported by TeraMetrics-based Ethernet and POS modules.

RFC 2460 defines six extension headers. Each is identified by a unique next header value. This appears in the previous header in the frame and indicates “the next header after this header” is going to be the identified type. The first extension header follows the IP header. There can be more than one extension header in the same packet, and certain extension headers can contain a variable number of fields, such as Type Length Value (TLV) values (hop-by-hop and destination options headers). Each extension header is an integer multiple of 8 octets long. This ensures the 8-octet alignment of subsequent headers.

RFC 2460 specifies a recommended order for IPv6 extension headers, as follows:

[IPv6 header]

Hop-by-hop options header
Destination options header (processed by first destination)
Routing header
Fragment header
Authentication header
Encapsulating security payload header
Destination options header (processed by final destination)

[Upper-layer protocol header: TCP, UDP]

The final extension header (the second destination options header) is followed by the upper-layer header (such as TCP).

With one exception, extension headers are not examined or processed by any node along a packet’s delivery path until the packet reaches the destination node. The exception is the hop-by-hop options header, which is examined by each node for the routing information it contains.

If a node does not recognize the next header value, the packet is discarded, and the ICMP parameter problem message is sent.

Important: SmartTracker will track an IPv6 frame with an extension header but will not track the portion of the frame that is the IPv6 extension header. When extension headers are present, tracking will only be up to the end of the IPv6 header. The extension headers and upper layer protocols will not be tracked properly.

Limitations When Configuring IPv6 Extension Headers

When using IPv6 extension headers, there are two items that you need to take into account:

- Aggregated header length versus header length limit
- Aggregated header length versus frame length.

The maximum length for the customizable header of a flow on the TeraMetrics modules is 128 bytes. Therefore, the number of extension headers that can be inserted into a flow's test frames is limited. The combined header length, including the extension headers, must not exceed 128 bytes. The available bytes for the IPv6 extension headers can be calculated as follows:

> Available bytes = 128 bytes - (layer 2 header size + IPv6 header size + upper layer protocol header size)

EXAMPLE FOR ETHERNET:

> Maximum header length = 128 bytes
> Layer 2 header = Ethernet = 14 bytes
> IPv6 header = 40 bytes
> Upper layer protocol = TCP = 20 bytes
>
> 128 - (14 + 40 + 20) = 128 - (74) = 54 bytes

Important: Inserting extension headers increases the minimum required size of a flow's test frames. Verify that the frame size specified for the flow is large enough to constitute a valid test frame.

The minimum frame length required can be calculated as follows:

> Minimum frame length required = layer 2 header + IPv6 header + IPv6 extension headers + upper layer protocol header + signature field + CRC

EXAMPLE FOR ETHERNET:

> Layer 2 header = Ethernet = 14 bytes
> IPv6 header = 40 bytes
> IPv6 extension header = fragment = 8 bytes
> Upper layer protocol = TCP = 20 bytes
> Signature field = 18 bytes
> CRC = CRC32 = 4 bytes
>
> 14 + 40 + 8 + 20 + 18 + 4 = 104 bytes

Note: A layer 2 header may include VLAN tags.

Setting Up IPv6 Extension Headers

To insert IPv6 extension headers into the test frames of a flow:

1. Ensure the IPv6 is enabled for the port. On the *Cards* tab, the *IPv6 Capable* checkbox for the port should be selected.
2. Set up the flow(s) by using the Group Wizard on the *Groups* tab or the *SmartFlows* tab. (Each IPv6 header group relates to a single flow.)

3 Click the **SmartFlows>Traffic** tab. Highlight the flow of interest, and verify that the *IPv6 flow* checkbox in the *Flow Properties* pane is selected.

4 Click the **IPv6 Ext Headers** tab (on the **SmartFlows** tab).

5 Use the *Add* button to move each desired IPv6 extension header from the *Options* column to the *List* column. (To add a header, highlight it and click the **Add** button.)

Use the *Remove* button to remove an IPv6 extension header from the *List* column. (To remove a header, highlight it and click the **Remove** button.)

6 To view the contents of a header, highlight it in the *List* of headers. The *Header Description* pane displays all pertinent parameters.

Tip: Use the *Up/Down* buttons to reorder the IPv6 headers in the *List* column. (To reposition a header, highlight it and click the **Up** or **Down** button.)

You can insert headers in any order, and you can insert as many headers as are possible for the frame length of the flow and the custom header length. Observe the following guidelines:

- The frame length for the flow can limit the number of headers that may be inserted. Frame length is defined by the *Frame Length (with CRC)* field on the *Traffic* tab.
- The total space available for header information (of all types) is 128 bytes.
- If a flow is associated with a custom frame sizes sequence, the minimum frame length of the sequence is used. You set up these sequences by using the *Custom Frame Sizes (per flow)* option on the *Test Iterations* tab.
- If VLAN or VLAN stacking is enabled for the flow, this reduces the space available for IPv6 extension headers. Each VLAN tag uses 4 bytes.

For a description of the frame length tables for IPv6 extension headers as well as the IPv6 extension header fields, refer to the SmartFlow online Help. Press **F1** over the *IPv6 Ext Headers* tab, or select **Help** from the menu bar.

Copying IPv6 Extension Headers to Multiple Flows

IPv6 extension header information can be copied from one flow to other flows.

When copying extension headers, all headers from the source flow are copied to the destination flows. If the source flow does not have any header, the destination flow loses any headers that it contains. In such a case, it is equivalent to removing headers from the destination flow.

To copy header information over multiple flows:

1 Highlight the source flow in the left pane on the **SmartFlows** tab.

2 Use the mouse and Ctrl or Shift keys to select additional flows in the pane.

3 Right-click and select **IPv6 Flow>Extension Header>Copy Down** from the pop-up menu. The IPv6 extension headers defined for the source flow are copied to the target flows.

Notes:

- The IPv6 extension header sequences are copied to the target flow only when there is enough room in the target flow's frame size for all the extension headers being copied down.
- If the first selected flow is not an IPv6 flow, then the copy-down operation is aborted. If a destination flow is not an IPv6 flow, that flow is skipped.
- If the IPv6 flow does not have a qualified TX port, that flow is skipped.

Configuring ATM Ports

If you are doing ATM testing and using ATM cards, use the *Cards* tab to set ATM-specific port parameters and the *ATM* tab to specify how you want the ATM virtual circuits configured on each ATM port. See *"Setting Up ATM Ports" on page 121* for more information.

ATM Basics

Asynchronous Transfer Mode (ATM) is a method of formatting, multiplexing, cross-connecting, transmitting, and switching information in 53-byte cells.

Virtual Circuits (VCs) are logical connections between end stations that are uniquely identified on a physical interface by a Virtual Channel Identifier (VCI) in the ATM cell header.

Virtual paths (VPs) are a bundle of VCs and are identified by a Virtual Path Identifier (VPI) in the ATM cell header.

Figure 5-12 shows the relationship of VPs to VCs.

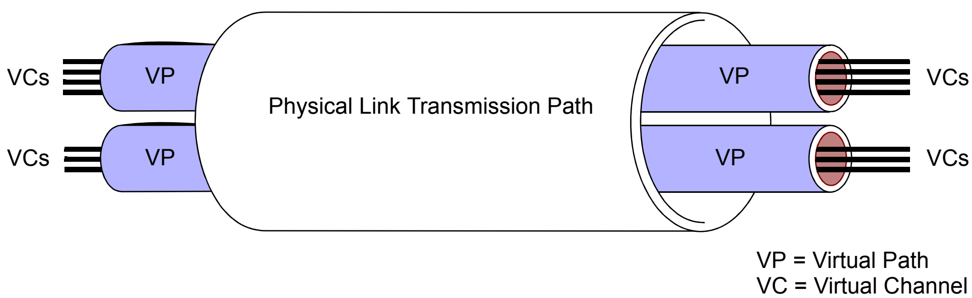

Figure 5-12. Virtual Paths and Virtual Circuits in ATM

ATM supports two types of circuits:

- Permanent Virtual Circuits (PVCs), which are always set up.
- Switched Virtual Circuits (SVCs), which are set up on demand.

ATM Cells

ATM sends fixed-length frames of 53 bytes: 5 bytes of header and 48 bytes of payload. There are two types of ATM cells:

- User-Network Interface (UNI) - used in cells sent by users.
- Network-to-Network Interface (NNI) - sent by switches to other switches.

ATM Support in SmartFlow

In SmartFlow, each virtual channel behaves like a separate port or sub-port. The virtual channel controls speed. All streams (flows) within a particular virtual channel have the same rate, assuming frame sizes are identical.

Some of SmartFlow's support for ATM includes:

- AAL5
- Up to 8191 PVCs (not SVCs)
- Up to 8191 VCs per port.

Note: Learning is not required.

SmartFlow supports ATM modules in SmartBits 6000x chassis.

ATM modules calculate and report latency the same way that TeraMetrics modules do. See *Table 12-2, "How Latency is Displayed, Depending on Flow's Receiving Port and Test Options," on page 353* for more information.

ATM modules do *not* support the following:

- IPv6
- Rates per flow
- Multicast traffic
- MPLS
- True random frame sizes
- True random test loads
- LAN emulation.

Note: ATM cards for the SmartBits 200/2000 are not supported in SmartFlow.

How SmartFlow Converts Frames to Cells

SmartFlow calculates and reports the rate for flows based on frames, not cells. ATM modules contain an AAL layer that segments frames into cells at the transmitting port and reassembles the cells back into frames at the receiving port.

Flows are transmitted with the number of cells needed to accommodate the frame size specified for the flow. AAL5 adds an 8-byte trailer to the upper layer protocol and padding to make the total divisible by 48. *Figure 5-13 on page 120* illustrates how this is done.

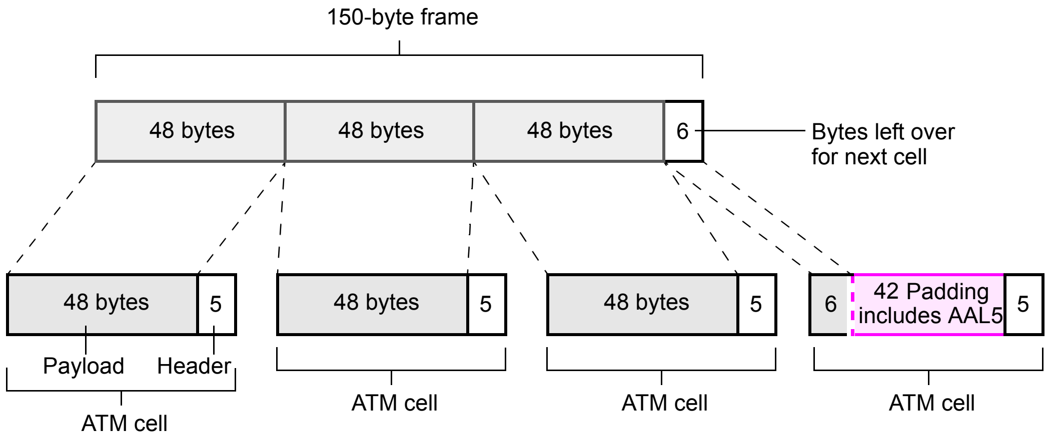

Figure 5-13. How SmartFlow Converts Frames to ATM Cells

The number of cells needed to insert into one IP frame is calculated with this formula:

1. Total length to be segmented =
LLC Routed IP header (8 bytes) + IP frame size without CRC + AAL5 trailer (8 bytes) + up to 47 bytes in padding
2. Number of Cells = (Total length to be segmented) / (53-5)

Note: If you use *VC Muxed Routed IP* for the encapsulation, the header is 0.

Example:
Assume an IP packet to be sent over an ATM network is 142 bytes without CRC. The VC used to send this packet has LLC Routed IP option specified for the encapsulation.

1. Total length to be segmented : 142 + 8 + 8 + 34 = 192
2. Number of cells : 192/48 = 4

Factors that Affect Speed in ATM

Each flow is transmitted in a round-robin fashion that affects the speed of each flow that is associated with the same transmitting port and VC.

You set the speed for a VC at the *ATM* tab. However the actual speed (rate) at which a VC transmits is affected by these factors:

- Port speed
- Total speed of all VCs on a port
(It must be less than or equal to the port speed.)
- Test load
(Only 100% yields the full speed specified for the VC.)

- Number of flows associated with the VC
- Frame size.

Example:
Let us say you set the speed of VC 0/32 on Port 1 to 64 Kbps. If you have five flows that each use VC 0/32 for the transmitting port, each flow is only transmitted at 1/5th of 64 Kbps if the test load is 100% and frame size is the same for each flow. If the frame size varies per flow, the bandwidth utilization of each flow varies accordingly, but the speed in terms of frames per second is still the same for each of the five flows.

Setting Up ATM Ports

Set up ATM ports at the *Cards* tab as you do with any type of port, with the addition of ATM-related fields. Then use the *ATM* tab to add and specify parameters for VCs for ATM ports. VCs must be defined on a port before the port can send or receive cells on that circuit. Each VC is equivalent to a subport and can be used as a source or destination port for flows.

The *ATM* tab displays all of the ATM ports that appear on the *Cards* tab. Each row on the ATM tab represents a VC associated with an ATM port. By default, each port has one VC. All cells transmitted on a VC have the header bits configured as specified at this tab. *Figure 5-14* shows an example of the *ATM* tab.

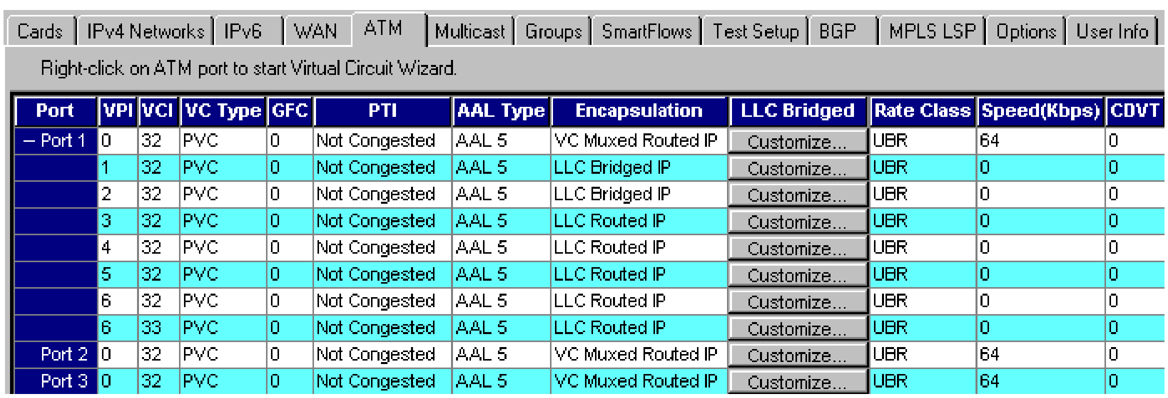

Port	VPI	VCI	VC Type	GFC	PTI	AAL Type	Encapsulation	LLC Bridged	Rate Class	Speed(Kbps)	CDVT
– Port 1	0	32	PVC	0	Not Congested	AAL 5	VC Muxed Routed IP	Customize...	UBR	64	0
	1	32	PVC	0	Not Congested	AAL 5	LLC Bridged IP	Customize...	UBR	0	0
	2	32	PVC	0	Not Congested	AAL 5	LLC Bridged IP	Customize...	UBR	0	0
	3	32	PVC	0	Not Congested	AAL 5	LLC Routed IP	Customize...	UBR	0	0
	4	32	PVC	0	Not Congested	AAL 5	LLC Routed IP	Customize...	UBR	0	0
	5	32	PVC	0	Not Congested	AAL 5	LLC Routed IP	Customize...	UBR	0	0
	6	32	PVC	0	Not Congested	AAL 5	LLC Routed IP	Customize...	UBR	0	0
	6	33	PVC	0	Not Congested	AAL 5	LLC Routed IP	Customize...	UBR	0	0
Port 2	0	32	PVC	0	Not Congested	AAL 5	VC Muxed Routed IP	Customize...	UBR	64	0
Port 3	0	32	PVC	0	Not Congested	AAL 5	VC Muxed Routed IP	Customize...	UBR	64	0

Figure 5-14. ATM Tab

As with the *Cards* tab, some columns in the *ATM* tab allow you to copy and paste cell values, copy down values, or change highlighted cells by incrementing, decrementing, or randomly filling them. Highlight the cell and right-click to access the menu.

The hierarchy of ATM entities is as follows:

- ATM port
- Virtual paths, which are associated with a port
- Virtual circuits, which are contained within virtual paths.

Each VC for a port is defined by a unique VPI/VCI pair of values. Idle cells are identified by a VPI/VCI of 0/0.

Notes:

- The maximum number of VCs allowed per port is 8192.
- You cannot delete a VC if it is used by any existing flows.

Tips:

- You must configure a VC for both the source and destination of a flow.
- When setting up VCs, make sure that the VPI/VCI number for a virtual circuit matches the VPI/VCI on the DUT to which it is connected.

For a description of each field on the *ATM* tab, refer to the SmartFlow online Help. Press **F1** over the tab, or select **Help** from the menu bar. The online Help also contains procedures and detailed usage information about adding virtual circuits with the *VC Wizard*. Idle cell parameter information is also included.

Configuring WAN Ports

If you are doing Frame Relay testing and using WAN cards, use the *WAN* tab to specify how you want the WAN ports configured. Setting up WAN ports follows this hierarchy of structure, as illustrated in *Figure 5-15*.

1. **Port-level parameters (DS3).** Each DS3 can contain up to 28 physical T1 lines.
2. **T1 parameters.** Each T1 line can contain up to 24 channels or DS0s.
3. **Channel parameters.** A channel is a logical part of a T1 line. A channel can consist of one or more DS0s. Each DS3 line can have up to 672 channels (24 channels X 28 T1 lines per DS3 line).
4. **DS0 parameters.** Assign one or multiple DS0s to one channel, but each channel belonging to a T1 line must have a unique set of DS0s.
5. **PVC (for Frame Relay) or PPP connection parameters.** You can assign up to 1022 PVCs to one or more channels defined for a single DS3 port. No more than 1022 PVCs can be assigned to any one DS3 port. If any LMI option is enabled, the maximum number of PVCs allowed per channel is 60.

Note: If you use the *Clear Channel* option for the port, a single channel is allocated all of the bandwidth.

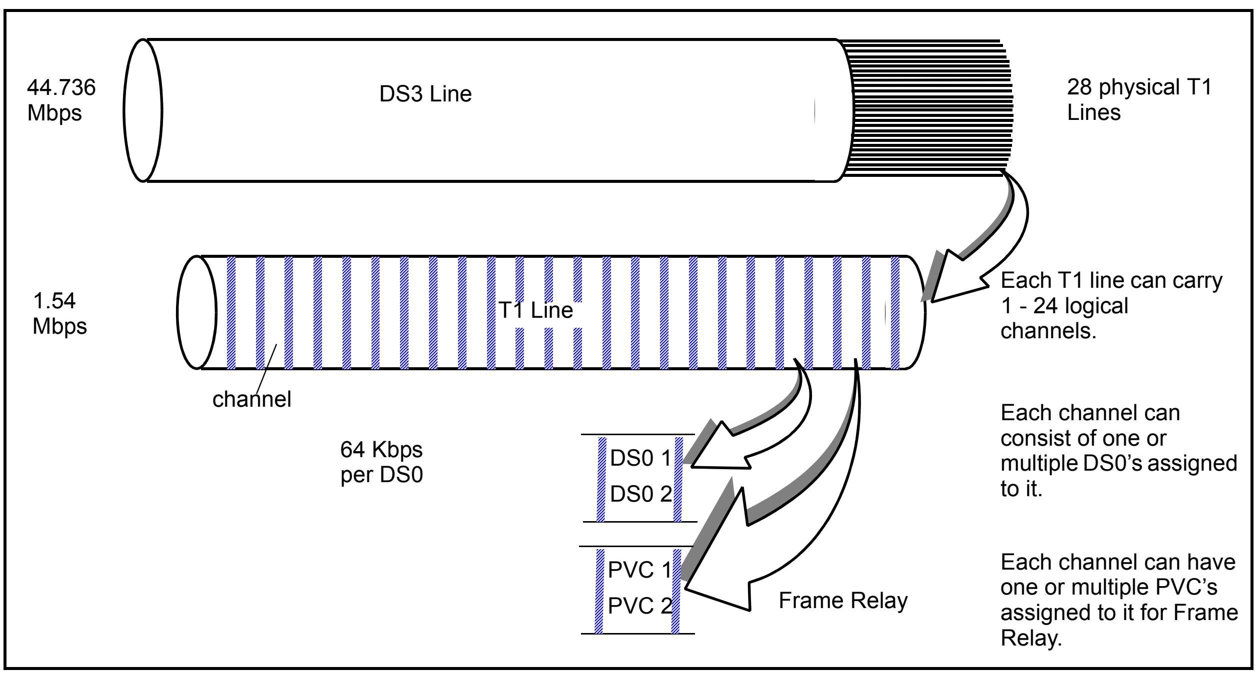

Figure 5-15. WAN Hierarchy of Structure

SmartFlow treats each PVC or PPP associated with a port as a separate port from which you can generate flows.

Using the WAN Tab

The *WAN* tab uses a tree structure to reflect the previously discussed hierarchy. *Figure 5-16* below shows what the *WAN* tab looks like with the tree collapsed.

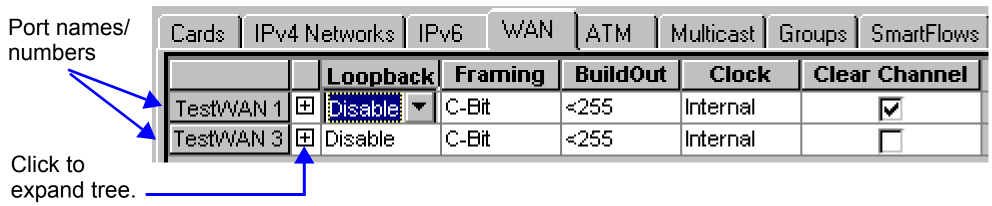

Figure 5-16. WAN Tab with Tree Collapsed

Notice that only port-level cells are displayed, such as for Loopback, Framing, and Buildout.

Figure 5-17 is what the *WAN* tab looks like with the tree expanded.

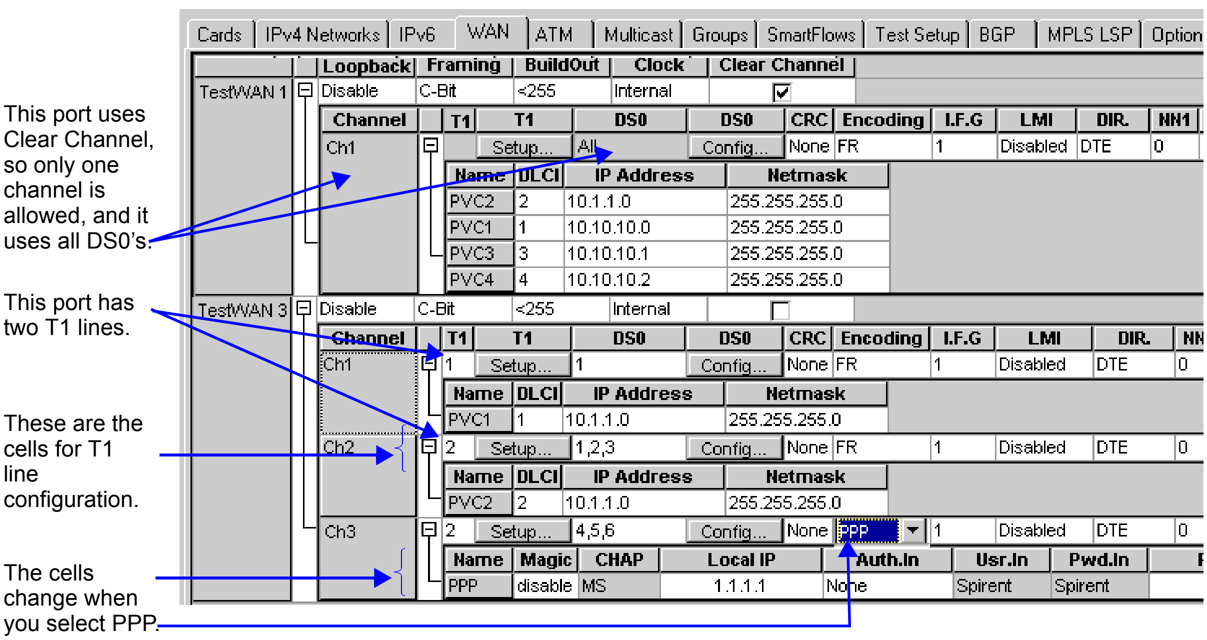

Figure 5-17. WAN Tab with Tree Expanded

In *Figure 5-17 on page 124*, because the first port uses clear channel, only one channel is allowed for the port and that channel uses all of the DS0s for a DS3 line. When clear channel is not selected (as with Port TestWAN 3), notice that a different set of fields display.

When you select PPP for the data link (encoding) method, the columns change to reflect this option.

For a description of each field on the *WAN* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar. The online Help also contains detailed procedures and information about how to use the *WAN* tab.

Guidelines for Setting Up WAN Ports

When setting up WAN ports, observe the following guidelines:

- The general port configuration for each card interface, such as framing, must match the configuration at the DUT interface.
- Clocking should be internal at one interface and loop-timed at the other, if ports are connected back-to-back. Otherwise, clocking should be set to be compatible with the DUT or SUT.
- DS0 selection also must match at the SmartCard interface and the DUT interface. However, on the two SmartCards, these configuration parameters can differ; they must only match the configuration of the terminating DUT interface.

Important: Make sure that each channel, DS0, and PVC that is added matches the DUT configuration.

Figure 5-18 on page 126 illustrates how configurations should match between cards and the DUT.

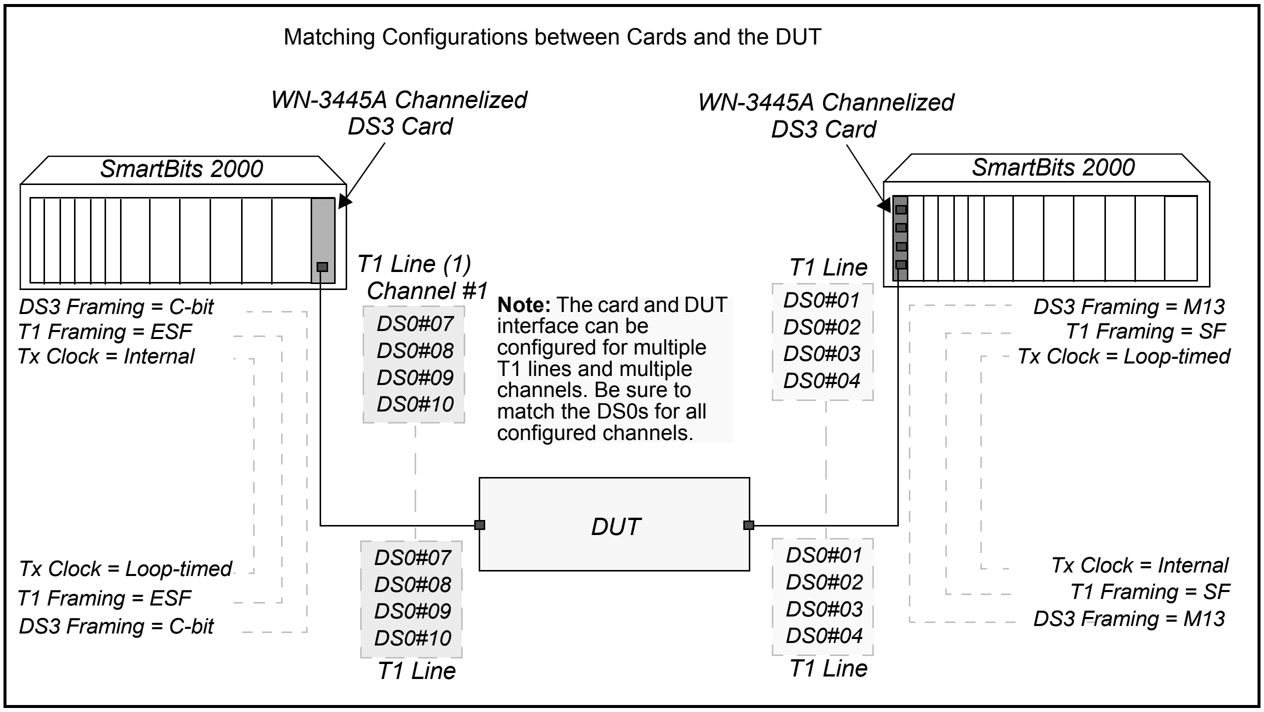

Figure 5-18. Matching Configurations between Cards and the DUT

The WAN port setup options vary according to what is selected for the link protocol (PPP or frame relay).

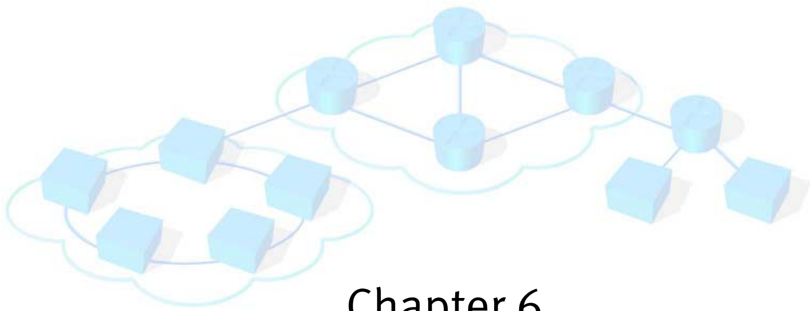

Chapter 6

Use the Wizard to Set up Flows/ Groups

This chapter provides information about how to plan flows and groups before setting them up. It explains when to use the Group Wizard, as opposed to manually setting up flows, and how to use the Group Wizard to automate flow and group generation. For the steps to set up flows without the Group Wizard, see *Chapter 7, “Work with Individual Flows.”*

In this chapter...

General Guidelines for Flow Setup and Testing

While every device is different and the way that you implement priority or VLANs is unique, these guidelines can help you effectively test whatever condition that you want to measure.

- Connect to chassis before you set up the ports, networks, and flows to ensure that flows and ports belong to the correct networks.
- Set up more groups and/or flows than the number of queues that exist in the DUT.
- Make sure to test *all* priorities that you have set up from *all* ports (instead of just one priority from one port). By testing all ports, you get a broad picture and in effect isolate and observe strictly priority issues, not port issues.
- For any priority implementation, separate traffic into high and low priority groups and/or flows.
- Use graph results for a quick visual indicator of what is occurring, but use the tabular data to analyze the situation.

See *"IPv6 Test Setup" on page 298* for guidelines on flow setup and testing for IPv6.

Planning the Flows and Groups

There are various ways to set up and differentiate flows and groups, such as by type of traffic or priority level. But before you set up any flows and groups in SmartFlow, know what characteristic or type of traffic that you want to see handled by the DUT. First answer the question: *What do you want to track?* The answer will be the basis for what flows and groups you set up in SmartFlow.

If you want to test a priority implementation, set up flows that mimic your implementation of priority. However, regardless of your priority implementation, you should separate the traffic by creating groups and flows according to this:

- High priority traffic
  Includes any critical traffic such as voice and video.
- Low priority traffic
  Includes stock quotes, news bulletins, and web updates.

Note: If a flow is not in a group, it is still transmitted if it is enabled. Flows can be members of multiple groups or no group at all.

Examples

You want to know: "*How does the device handle traffic from high to low priorities?*"

Set up flows with priorities that range from high to low. Then set up a group containing flows with only high priorities, and another group containing flows with only low priority.

You want to know: "*How does the device handle Telnet traffic versus email traffic?*"

Set up some flows with Telnet as the port type and other flows with SMTP for the port type.

Why Group Flows?

SmartFlow uses groups for the purpose of displaying test results. If a number of flows are similar or in some way share a common characteristic that you want to track (such as VLAN, protocol, or priority), then include them in a group to more easily track and view that characteristic in test results. You can also include the same flow in more than one group for comparison. The key to setting up a group is knowing what you want to track.

Examples

You want to track: *aggregate performance of Telnet traffic (TCP port 23) through a device.*

Group all of the Telnet flows together.

You want to set up: *some flows with high priority, some with low priority, plus some with TCP and some with UDP protocols.*

Set up a high priority group, a low priority group, a TCP group, and a UDP group.

Selecting Groups for Your Test

Once you have created groups of flows, you can select specific groups to be included in your test, while excluding others. Use the checkbox next to each group name to do this.

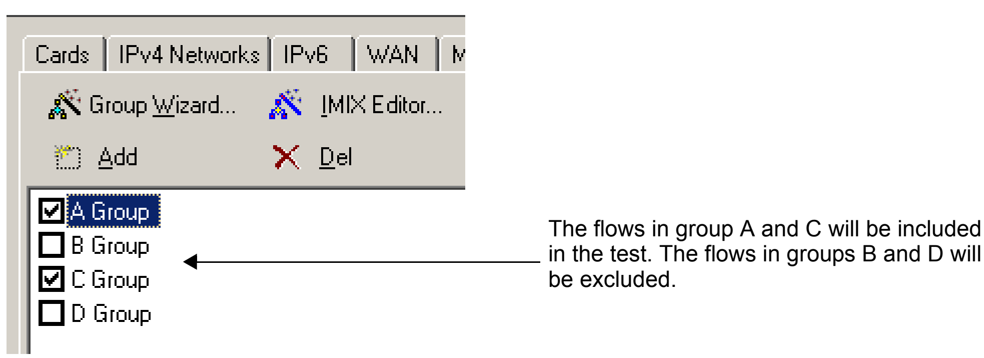

When the test runs, port and flows that belong to the selected groups are included in the test. Individual flows not associated with any group are also included. If a flow is included in multiple groups, it is included if any of the groups is selected.

Your group selections are stored in the .flo configuration file.

Sample Flow/Group Setup

Let us say you want to see how the device handles protocols as well as priority. Specifically, you want to see how the device handles:

- IP, TCP, and UDP protocols
- High, low, and medium priority traffic.

Assume that you are using three test ports in a mesh test. Set up the following groups:

- Three groups for each protocol
- Three groups for each level of priority.

Figure 6-1 illustrates this. With this configuration of flows and groups, you can observe if any one flow is getting precedence over another and what port each flow came from.

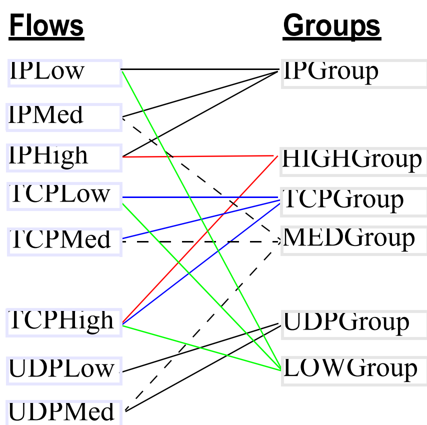

Figure 6-1. Example Flow/Group Configuration to Test Priority

Notice that each flow belongs to two groups. For example, the IPMed flow belongs to both the IPGroup and the MEDGroup groups.

If you used the default TOS priority settings of 0, 3 and 7 (which roughly represent low, medium and high priority), the SmartFlow Group Wizard generates IP flows as shown in *Figure 6-2 on page 131*.

Note: To track how the device handles priority traffic, configure two or more ports to overload one of the ports.

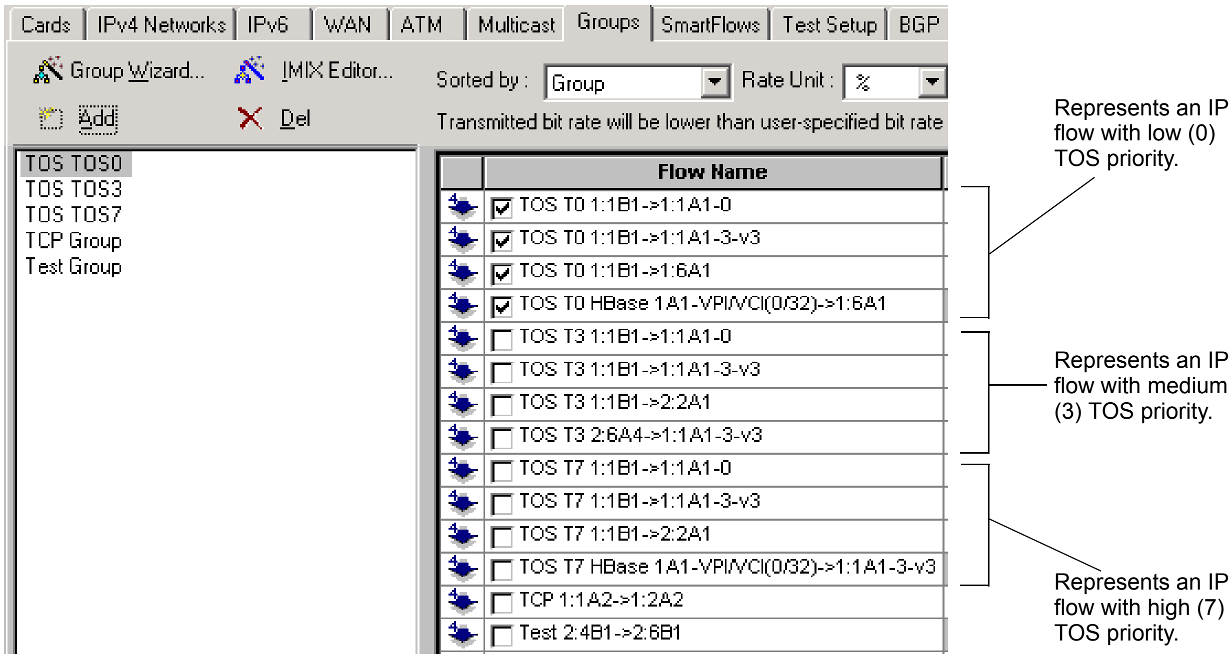

Figure 6-2. Resulting Flows from Example Priority Configuration

Since prioritization is not visible until there is congestion, the load and test results for the three ports may look something like this:

TX Line Rate (Mbps)	Total % of Line Rate DUT is Receiving from All Ports (Mbps)
10	30
20	60
30	90
	Packet loss and prioritization starts here.
40	120
50	150
60	180
70	210
80	240
90	270
100	300

About the Group Wizard

Use the Group Wizard as your standard method of creating groups and flows. The Group Wizard sets up multiple flows and groups all at the same time. Each flow in the group shares the same IP protocol and VLAN, TOS/IP Precedence, or Diffserv settings.

Use the Group Wizard if:

- You want a fast and easy way to automatically create flows and groups.
- You have a lot of flows and groups to set up.
- You want a mesh test with automatically balanced traffic on all ports.
- The device supports priority schemes and you want to vary groups only by TOS, DiffServ, and VLAN priority.
- You want to easily configure traffic based on multiple networks and/or VLANs per port.
- You want to automatically assign rates per flow according to VLAN and IP priority while simultaneously creating flows and groups.

How the Wizard Sets Up Flows and Groups

The Group Wizard allows you to quickly set up multiple IP flows and groups with:

- Automatically varied or uniform TOS, Diffserv, or VLAN priority settings.

  Based on your traffic pattern, the Group Wizard automatically generates a group and set of flows for each priority setting in the selected routing scheme.
- The *same* protocol type (with or without priority settings).

  Based on your traffic pattern, the Group Wizard automatically generates a group and set of flows for the protocol that you select.
- Randomly generated or the same port types (for TCP or UDP only).

  Based on your traffic pattern, the Group Wizard automatically generates a group and set of flows for each. The Group Wizard can automatically vary the port type or by one specific type.
- Transmitters in the waterfall pattern (for waterfall traffic pattern only).

  Based on the waterfall pattern, the Group Wizard automatically generates a group and set of flows based on the transmitters in the pattern. You can also use this together with the priority option.

Once you generate the flows, you can view the attributes of the individual flows. You can also modify the individual flows. See *"Modifying Flows and Groups Created with the Wizard" on page 169*.

For information on setting up flows and groups individually, see *Chapter 7, "Work with Individual Flows."*

How the Wizard Balances Traffic in a Mesh Pattern

When you use the Group Wizard for a mesh traffic pattern (every port in the group sending to each other), it automatically balances traffic load.

If you manually set up a group with a mesh traffic pattern, the order in which you set up the flows is important if you want each port to receive an equal load. If you set up flows in the wrong order, it can result in overloading on one port over the others, and test results will not be accurate. (In addition, a port can only transmit one flow at one instant in time.)

What is a Balanced Traffic Load?

The diagram of a fully-meshed test in *Figure 6-3* shows how each port sends to all other ports:

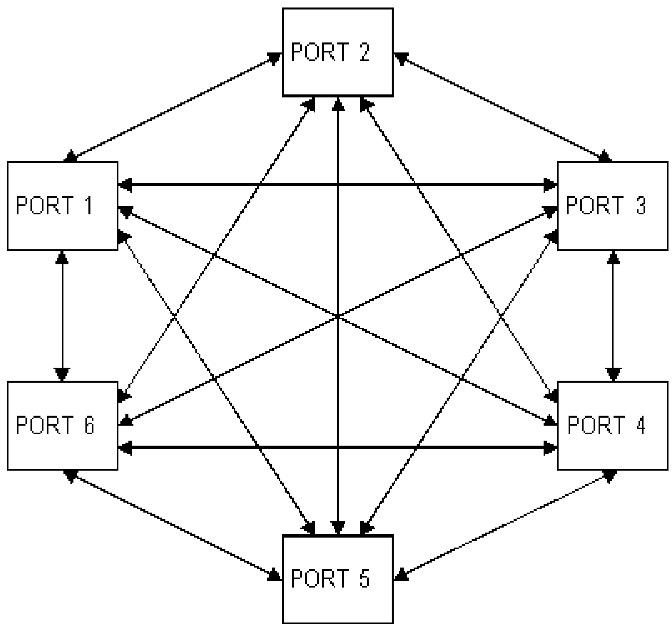

Figure 6-3. All Ports Sent to All Other Ports in a Fully-meshed Test

Each port in the test sends frames to all other ports in a round-robin fashion.

Table 6-1 on page 134 shows how each port in the test transmits frames to all other ports in the test. In this example, there are six ports with one address per port.

Table 6-1. Example of Transmission Order in Fully-Meshed Test

Source Port	Destination Ports (in Order of Transmission)					
Port #1	2	3	4	5	6	2 …
Port #2	3	4	5	6	1	3 …
Port #3	4	5	6	1	2	4 …
Port #4	5	6	1	2	3	5 …
Port #5	6	1	2	3	4	6 …
Port #6	1	2	3	4	5	1 …

There is an equal distribution of destination addresses to keep the test balanced. All ports are equally loaded during the test so that no one port is overloaded.

Important: If you manually set up flows, to keep the load balanced in a mesh traffic pattern, make sure that each flow contains an equal number of addresses per port.

How Does Traffic Become Unbalanced?

The following example demonstrates how one port in a mesh pattern can end up receiving traffic from two ports at once if the flow order is incorrect:

Let us say you want to set up a group with three ports (each are full duplex) in a mesh pattern, with Ports 1, 2, and 3.

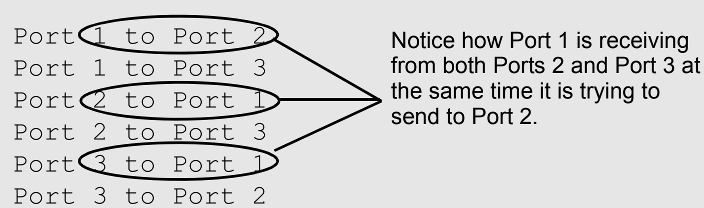

How the Wizard Balances Traffic in a Backbone Pattern

When you use the Group Wizard for a bi-directional backbone traffic pattern (every port transmitting to whatever port it is receiving from), it automatically balances traffic load per the IETF draft *Benchmarking Methodology for LAN Switching Devices*.

If you manually set up a group with a bi-directional backbone traffic pattern, the order in which you set up the flows is important if you want each port to receive an equal load. If you set up flows in the wrong order, it can result in overloading on one port over the others, and test results will not be accurate. (In addition, a port can only transmit one flow at one instant in time.)

The diagram of a bi-directional backbone test in *Figure 6-4* shows how each port sends to the other ports:

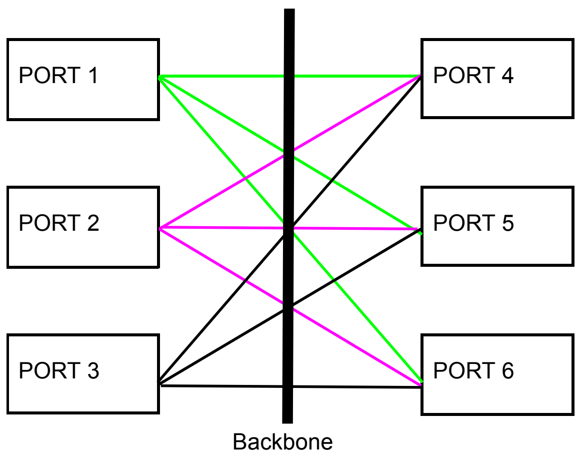

Figure 6-4. How Each Port Sends to Other Ports in a Bi-directional Backbone Test

Each port in the test sends frames to all other ports in a round-robin fashion.

Table 6-2 on page 131 shows how each port in the test transmits frames to all other ports in the test. In this example, there are six ports (three on each side of the backbone) with one address per port.

Table 6-2. Example of Transmission Order in Bi-directional Backbone Test

Source Port	Destination Ports (in Order of Transmission)		
Port #1	4	5	6
Port #2	5	6	4
Port #3	6	4	5
Port #4	1	2	3
Port #5	2	3	1
Port #6	3	1	2

There is an equal distribution of destination addresses to keep the test balanced. All ports are equally loaded during the test so that no one port is overloaded.

Important: If you manually set up flows, to keep the load balanced in a bi-directional backbone traffic pattern, make sure that each flow contains an equal number of addresses per port.

Understanding Group Wizard Traffic Patterns

The Group Wizard allows you to configure traffic in four possible patterns:

- **Backbone** (one-to-many or many-to-one in either direction or bi-directional)
- **Fully-meshed** (every port transmitting to every port)
- **Pair** (one or more pairs of ports with transmission within each pair in either direction or bi-directional)
- **Waterfall** (all ports transmit and receive in a many-to-one, partial mesh).

These traffic patterns and their resulting flows are best understood with examples shown in the sections that follow.

Note: The Group Wizard traffic patterns for multicast traffic differ somewhat in configuration from unicast traffic. The *Waterfall* pattern is not available for multicast traffic. For more information about Group Wizard patterns with multicast traffic, see *"Group Wizard Traffic Patterns for Multicast" on page 489*.

Backbone Traffic Pattern

Based on the uni-directional backbone pattern diagram below and no priority, the Group Wizard would generate one group containing flows 1-2, 1-3, and 1-4.

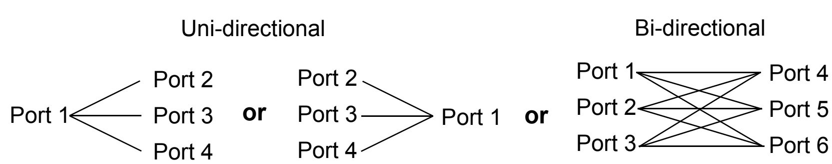

Let us say you select three ports in a many-to-one traffic configuration. Traffic is unidirectional, with ports in column A transmitting to the ports in column B, as shown in *Figure 6-5 on page 138*.

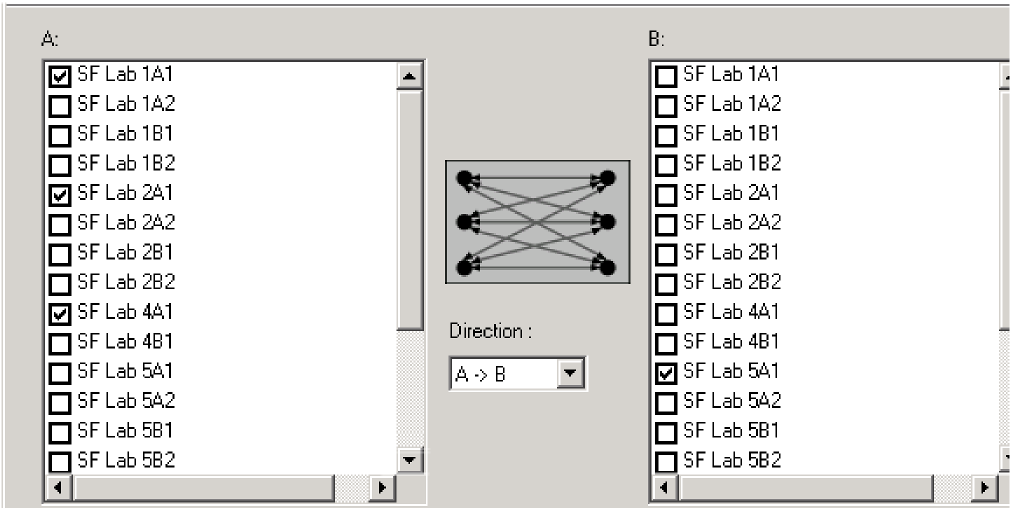

Figure 6-5. Example of Backbone Many-to-One Traffic Pattern

If you also chose to generate multiple groups by TOS precedence (at various levels), the SmartFlow Group Wizard generates the following flows from a backbone many-to-one, uni-directional traffic pattern:

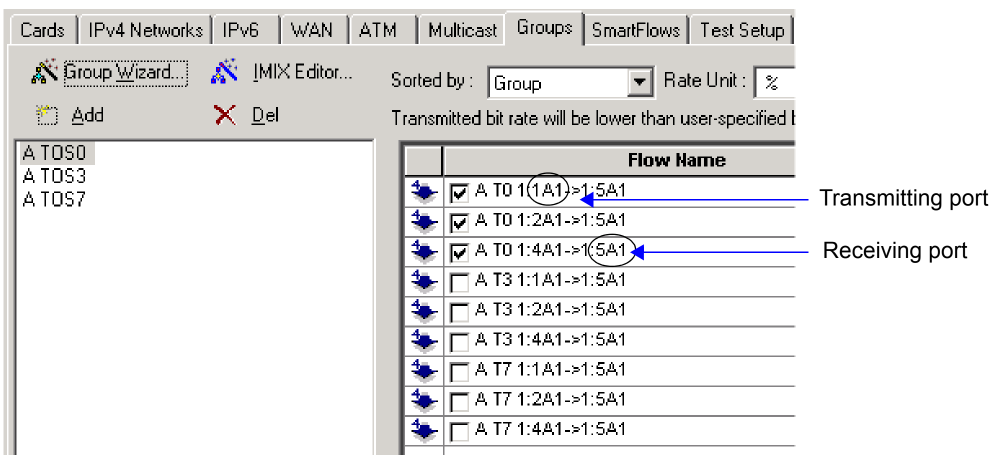

Notice that for each flow in each group, the transmitting port (in the *Flow Name* column on the right) is transmitting to port 4, the receiving port. This is because the traffic direction is unidirectional in the direction from A to B.

If the traffic direction was reversed, with the port in column B transmitting to the ports in column A, the flows created would look like this:

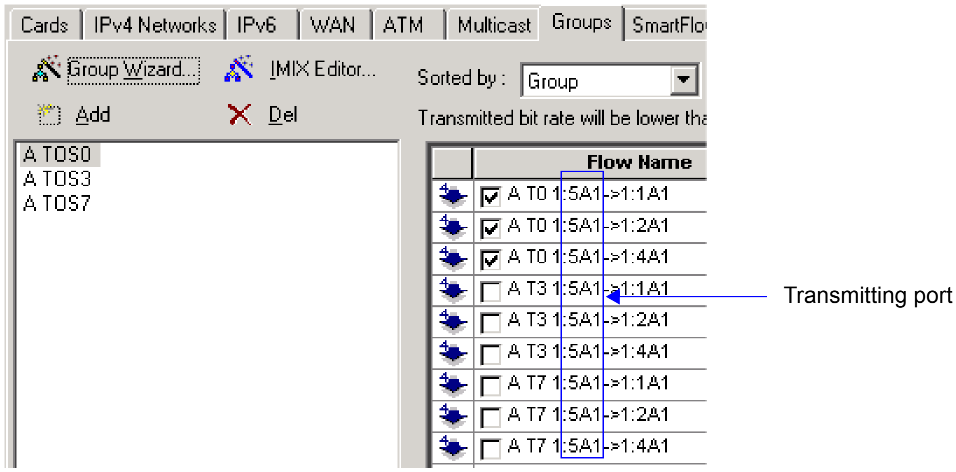

Notice that for each flow in each group, the transmitting port (from the column on the right) is port 4. This is because the traffic direction is uni-directional in the direction from B to A.

Bi-directional Backbone Pattern

Let us say you use the same port setup as the unidirectional traffic (*Figure 6-5 on page 138*), but now select bi-directional traffic (A to B and B to A). If you also chose to generate multiple groups by TOS precedence, the flows generated look like this:

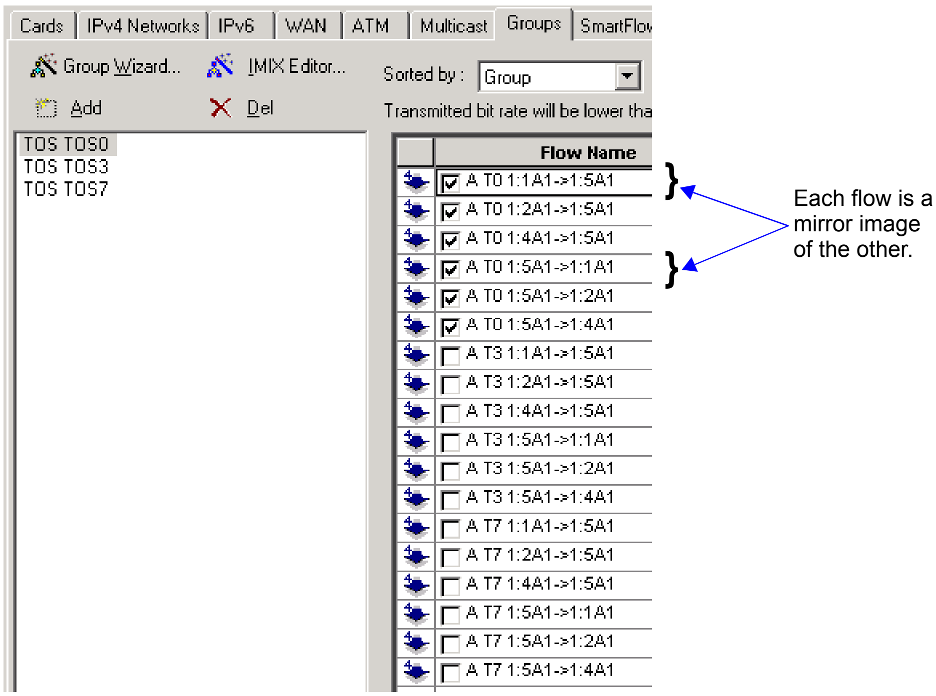

Notice that when you choose a bi-directional backbone traffic pattern, the Group Wizard creates twice as many flows. Each flow is a mirror image of the other: If port 1 transmits to port 2, port 2 also transmits to port 1.

If you select bi-directional for the traffic direction, the Group Wizard automatically balances the traffic load by transmitting in a round-robin fashion so that no port is overloaded at any time. For more information about load balancing, see *"How the Wizard Balances Traffic in a Backbone Pattern" on page 135*.

Fully-meshed Traffic Pattern

Based on the diagram below, with a fully-meshed configuration and no priority, the Group Wizard would generate one group that includes flows with every combination among the selected ports (1-2, 1-3, 2-1, 2-3, 3-1, 3-2).

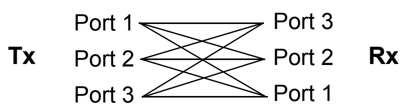

The *Fully-meshed* traffic pattern consists of every port transmitting to every port. Each port transmits to every other port except itself. The *Wizard* generates the number of flows for a *Fully-meshed* traffic pattern using this formula, where *n* is the number of ports involved in the test: ***n* x (*n*-1).**

The Group Wizard automatically balances the traffic load by transmitting in a round-robin fashion so that no one port is overloaded at any time. For more information about load balancing, see *"How the Wizard Balances Traffic in a Mesh Pattern" on page 133*.

Figure 6-6 shows an example setup for a fully-meshed traffic pattern.

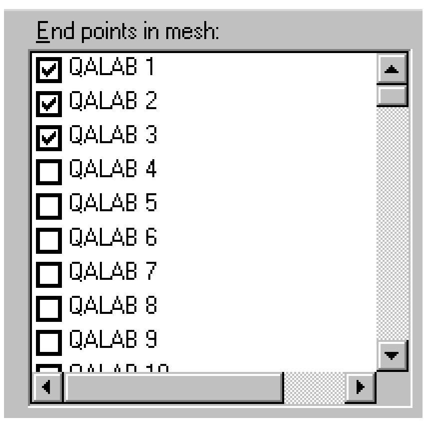

Figure 6-6. Example Setup for Fully Meshed Traffic Pattern

Based on this setup, if you also chose to generate multiple groups by TOS precedence, the Group Wizard creates these flows:

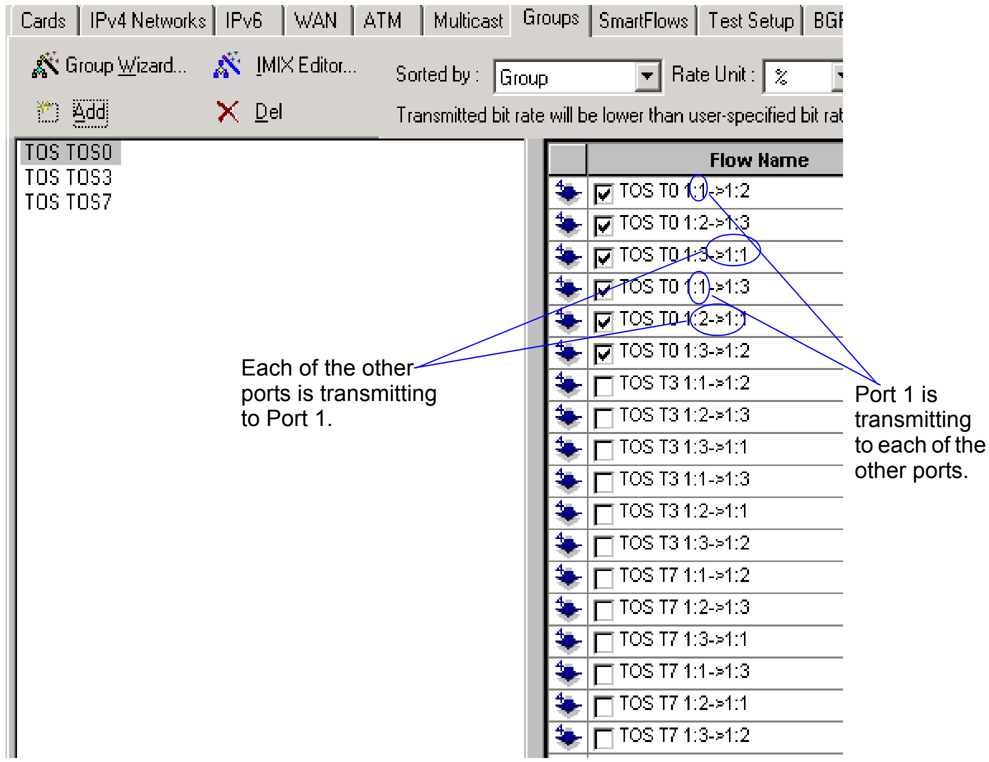

In the *Fully-meshed* traffic pattern, each port transmits to each port. The number of flows generated is `n x (n-1)`.

Pair Traffic Pattern

Based on the diagram above, with a pair pattern and no priority, the Group Wizard generates one group containing flows 1-2, 3-4, and 6-8.

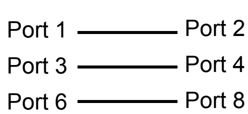

The *Pair* traffic pattern consists of one or more pairs of ports with transmission only within each pair (uni-directional or bi-directional).

Figure 6-7 on page 143 shows an example of a *Pair* traffic pattern setup consisting of two port pairs. Traffic is unidirectional, with the ports in column A transmitting to the ports in column B.

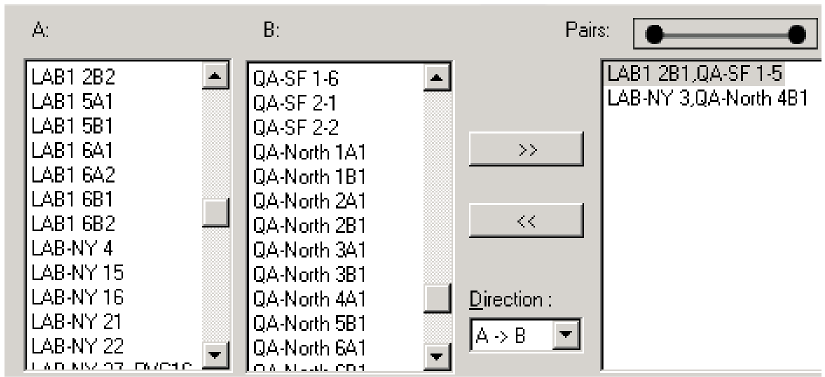

Figure 6-7. Example of Pair Traffic Pattern Setup

Based on the setup shown in *Figure 6-7*, if you also chose to generate multiple groups by TOS precedence, the *Wizard* creates these flows:

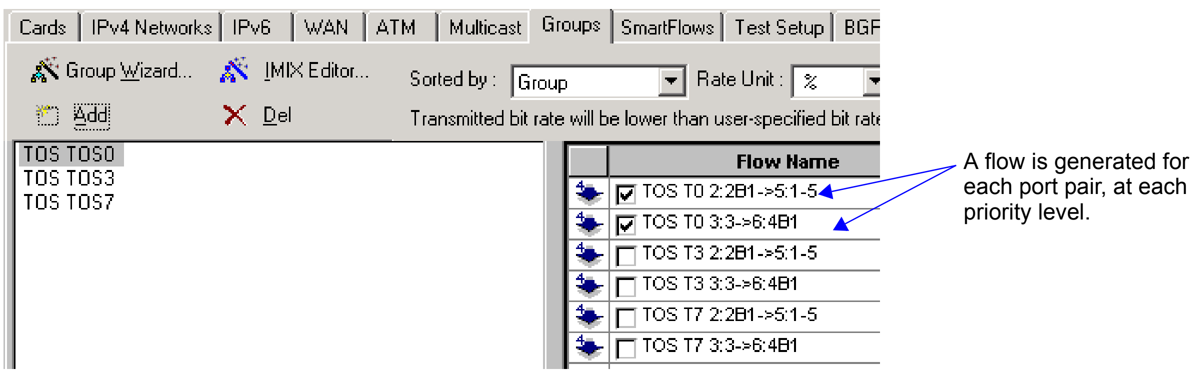

Tip: You can have a one-port configuration (same port for source and destination, such as for loopback testing). Once you select the *Pair* pattern, check the same port in list A and list B in the Group Wizard. Then select the **One port configuration** checkbox and move the “pair” into the *Pairs* list.

Bi-directional Pair Pattern

Let us say you use the same port setup as the unidirectional traffic in *Figure 6-7 on page 143*, but now select bi-directional traffic (A to B and B to A). If you also choose to generate multiple groups by TOS precedence, the flows generated look like this:

When you choose a bi-directional pair traffic pattern, twice the number of flows are created.

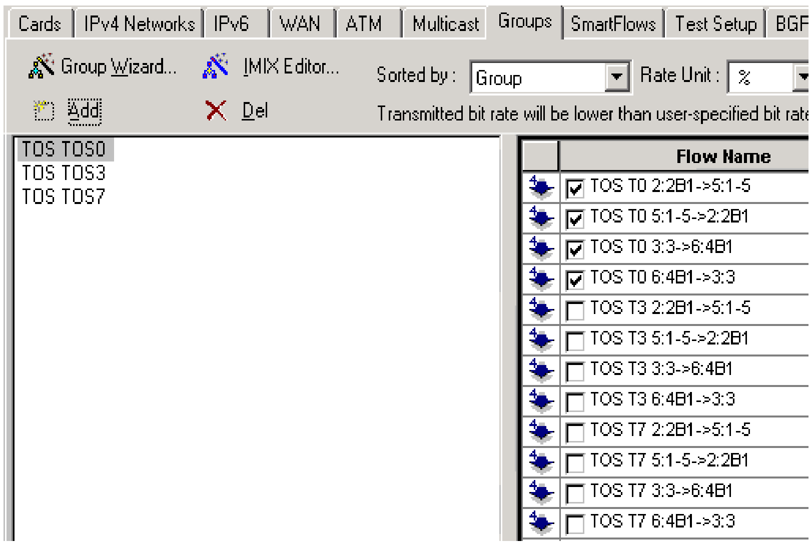

Waterfall Traffic Pattern

In the *Waterfall* traffic pattern, ports both transmit and receive in a many-to-one, partial mesh pattern. (The first set of transmitting ports only transmit, and the last group of receiving ports only receive.) It results in fewer flows than full or partial mesh tests because you can specify the number of transmitting and receiving ports in the pattern.

The *Waterfall* pattern is similar to a relay in that the receiving ports in the previous group become the transmitting ports in the next group. The receiving ports in the next group are drawn (in port number order) from the remaining unused ports in the range of ports that you selected. This continues until these two conditions exist:

- There are no unused ports remaining from the ports selected to be transmitters.
- There is at least one receiving port. (The interleave factor must be less than or equal to the number of transmitters).

The interleave factor is one of the criteria that you set for the waterfall pattern. It determines the number of ports that both transmit and receive in each group.

Note: The *Waterfall* pattern does not utilize cyclic flows or multicast flows. For more information about cyclic flows, see *"What is a Cyclic SmartFlow?" on page 31*. For more information about multicast traffic, see *Chapter 19, "IP Multicast Testing."*

The *Waterfall* pattern is more flexible than the other traffic patterns, with more variations possible. Read the sections that follow to understand how to use this powerful option. Refer to the online Help for a procedure about how to specify *Waterfall* pattern criteria in the Group Wizard.

Why Use the Waterfall Pattern

Existing traffic patterns in the Group Wizard (*Backbone, Fully-meshed, Pair*) can easily result in a very large number of flows, which can be somewhat unwieldy and overload a DUT. You may not need this many flows for your testing purposes. This situation occurs particularly if you select a fully-meshed or a partially-meshed pattern such as backbone. For example, if you use 10 ports in a fully-meshed test, it results in at least 90 flows. If you select to generate multiple groups (based on priority levels), the number of flows increases even more.

The waterfall pattern allows you to:

- Send flows to and from ports with a limit to the number of resulting flows.
- Oversubscribe ports.
- Have all ports active at the same time.
- Create less flows.

When to Use the Waterfall Pattern

Use the waterfall pattern if you want to:

- Limit the number of flows generated, yet be able to oversubscribe ports.
  The *Fully-meshed* pattern may generate more flows than needed, but the *Pair* pattern may not generate enough. The waterfall architecture of more transmitters than receivers causes congestion on the ports and allows you to test priority schemes.
- Test QoS.
  You can use the *Generate multiple groups by* option in the *Group Wizard - Characteristics* page to generate multiple groups by type of service together with the waterfall pattern. This results in separate waterfall traffic patterns in separate groups for each combination of priority bits selected using the *Variables* button.
- Automatically create multiple, partially-meshed (backbone) groups at one time. (The number of groups created is a factor of the traffic pattern criteria.)
- Simulate various oversubscription scenarios.
- Test individual modules of a DUT. (Some DUTs contain modules with multiple ports.) Use an interleave factor of 0 to keep all of the traffic within one module.
- Test multiple modules of a DUT.
  Increase the interleave factor from zero to correspond to the number of modules that you want to test.

Waterfall Examples with Various Interleave Factors

The *Waterfall* traffic pattern is easiest to understand with examples and diagrams. These examples show patterns whose criteria does *not* include grouping by transmitter or priority. For information about these, see sections *"Generating Multiple Groups by Transmitter (No Priority)" on page 159* and *"Generating Multiple Waterfall Groups by Priority Only" on page 163*.

Example: Interleave Factor = 0

Ports Selected: **1- 8**
Number of Transmitters: **2**
Number of Receivers: **2**
Interleave factor: **0**

Figure 6-8 shows the pattern of flows that is created with this set of criteria.

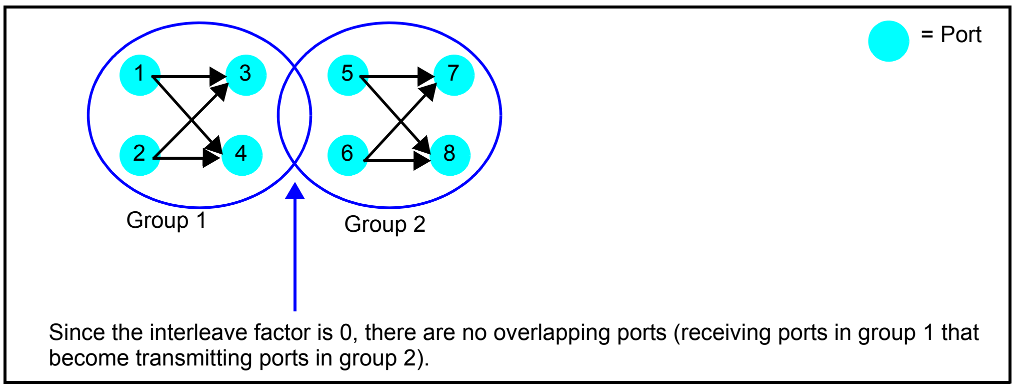

Figure 6-8. Waterfall Pattern with Interleave Factor of 0

Example: Interleave Factor = 2

Ports Selected: **1- 8**
Number of Transmitters: **2**
Number of Receivers: **2**
Interleave factor: **2**

Figure 6-9 shows the pattern of flows that is created with this set of criteria.

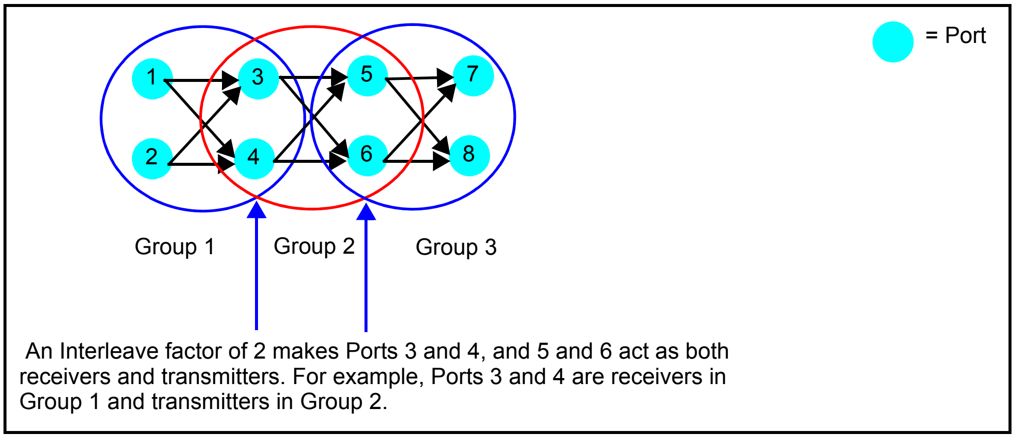

Figure 6-9. Waterfall Pattern with Interleave Factor of 2

With an interleave factor of 2, two receiving ports in each group become transmitting ports in the next group.

You can "close the link" in this configuration by manually setting up a group where ports 7 and 8 transmit to ports 1 and 2.

Example: Interleave Factor = 3

Ports Selected: **1- 16**
Number of Transmitters: **4**
Number of Receivers: **3**
Interleave factor: **3**

Figure 6-10 shows the pattern of flows that is created with this set of criteria.

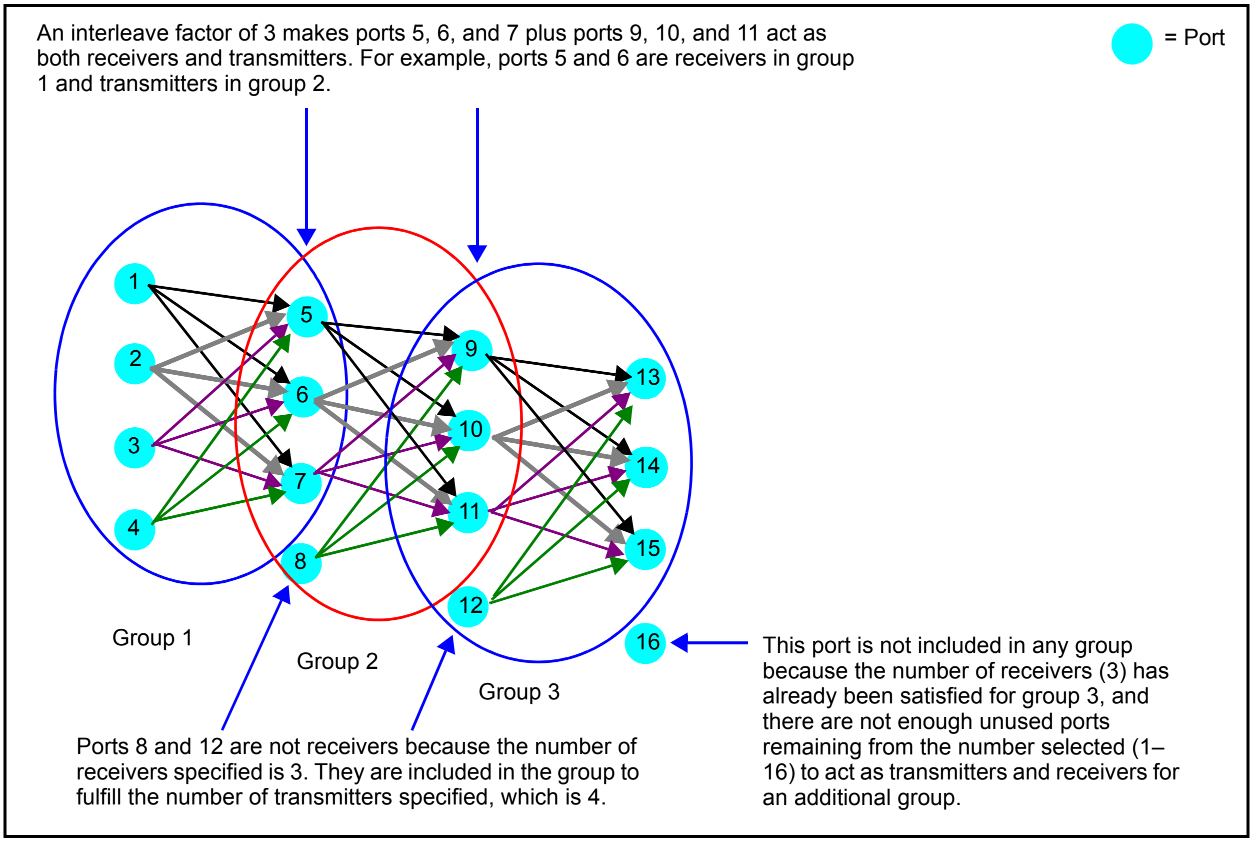

Figure 6-10. Waterfall Pattern with Interleave Factor of 3

With an interleave factor of 3, three receiving ports in each group become transmitting ports in the next group. In this example, port 8 does not receive any traffic in group 1 because you have four transmitters but only three receivers (and port 8 would be the fourth receiving port, which exceeds the specified number of transmitters).

Using the Group Wizard

The Group Wizard consists of four pages:

- Group Wizard - Traffic Selection (first)
- Group Wizard - Traffic Configuration (second)
- Group Wizard - Characteristics (third)
- Group Wizard - Multiple Flows (fourth).

Note: Since the *Waterfall* pattern does not support cyclic flows, the Multiple Flows page does not appear with this pattern.

If you plan on creating flows that use a network or VLAN added to the port as the flow's source or destination but also want to keep the port's original base MAC address the same, you must add a line to the `SmrtFlow.ini` file. (See the online Help for information about how to keep the same MAC address with multiple networks on a port.)

To access the Group Wizard, click the **Groups** tab and then the **Group Wizard** button. The *Group Wizard - Traffic Selection* page opens.

Defining Traffic Type and Pattern in the Wizard

To access the Group Wizard, click the **Groups** tab and then click the **Group Wizard** button. The *Group Wizard - Traffic Selection* page (*Figure 6-11*) is the first of four pages. Use this page to define the traffic type and pattern for ports in the test. (Refer to *"Understanding Group Wizard Traffic Patterns" on page 137*).

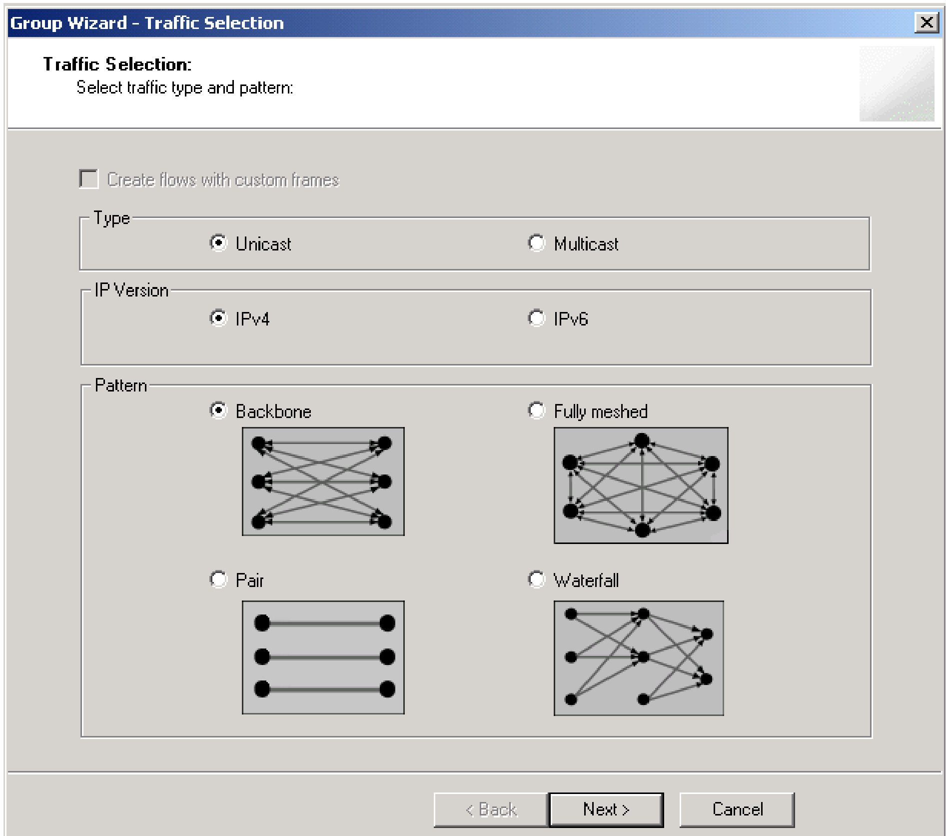

Figure 6-11. Group Wizard - Traffic Selection Page to Select Traffic Type and Pattern

The *Create flows with custom frames* box is disabled unless you have first created some custom frames at the *Test Setup>Test Iterations* tab.

Note: The *Waterfall* traffic pattern setup and behavior is different from the other traffic patterns. Refer to the online Help for information about how to specify *Waterfall* pattern criteria in the Group Wizard.

For detailed information about the *Traffic Selection* page, refer to the SmartFlow online Help. Press **F1** over this page or select **Help** from the menu bar. The Help contains detailed information and procedures for all pages in the Group Wizard.

Now select the ports using the second page, the *Group Wizard - Traffic Configuration* page.

Selecting Ports in the Wizard

After you select the traffic type and pattern, use the *Group Wizard - Traffic Configuration* page (*Figure 6-12*) to define the ports and traffic direction for the test. This is the second of four pages in the Group Wizard. The *Mesh* pattern provides some traffic balancing options.

For a detailed information on the *Traffic Configuration* page, refer to the SmartFlow online Help. Press **F1** over this page, or select **Help** from the menu bar. The online Help contains detailed information and procedures for all pages in the Group Wizard.

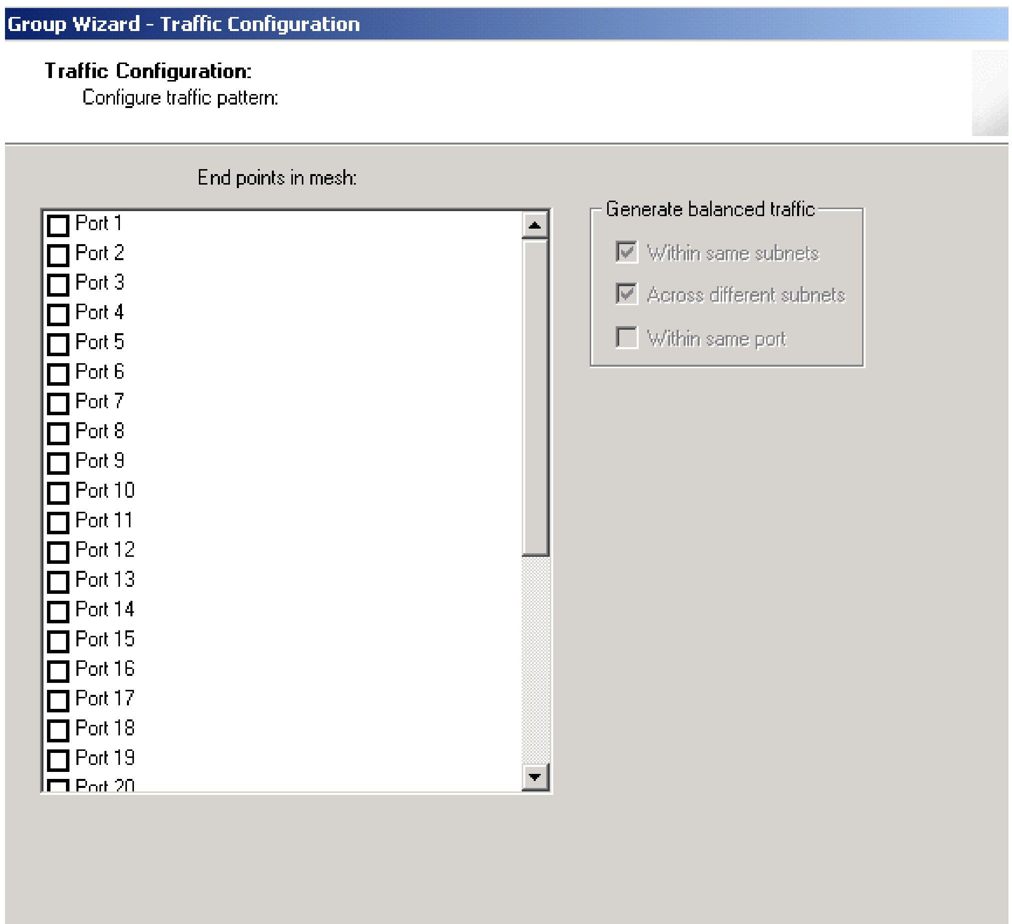

Figure 6-12. Group Wizard - Traffic Configuration Page Using Mesh Pattern

Defining Flow and Group Characteristics in the Wizard

After you have selected a traffic pattern and ports, use the *Group Wizard - Characteristics* page (*Figure 6-13*) to define the characteristics of flows and groups involved in the test. This is the third page of the Group Wizard.

For a detailed information on the *Group Wizard - Characteristics* page including the Diffserv class dialog box, refer to the SmartFlow online Help. Press **F1** over this page, or select **Help** from the menu bar. The online Help contains detailed information and procedures for all pages in the Group Wizard.

Note: When you run a test, *all* flows listed on the *Groups tab* (or *SmartFlows* tab) are sent. If you do not want to include a group, you must delete it before you launch the test. If you do not want flows sent in a test, disable the flow by unchecking it at the *SmartFlows* tab. See *"Including and Excluding Flows for a Test" on page 210* for more information.

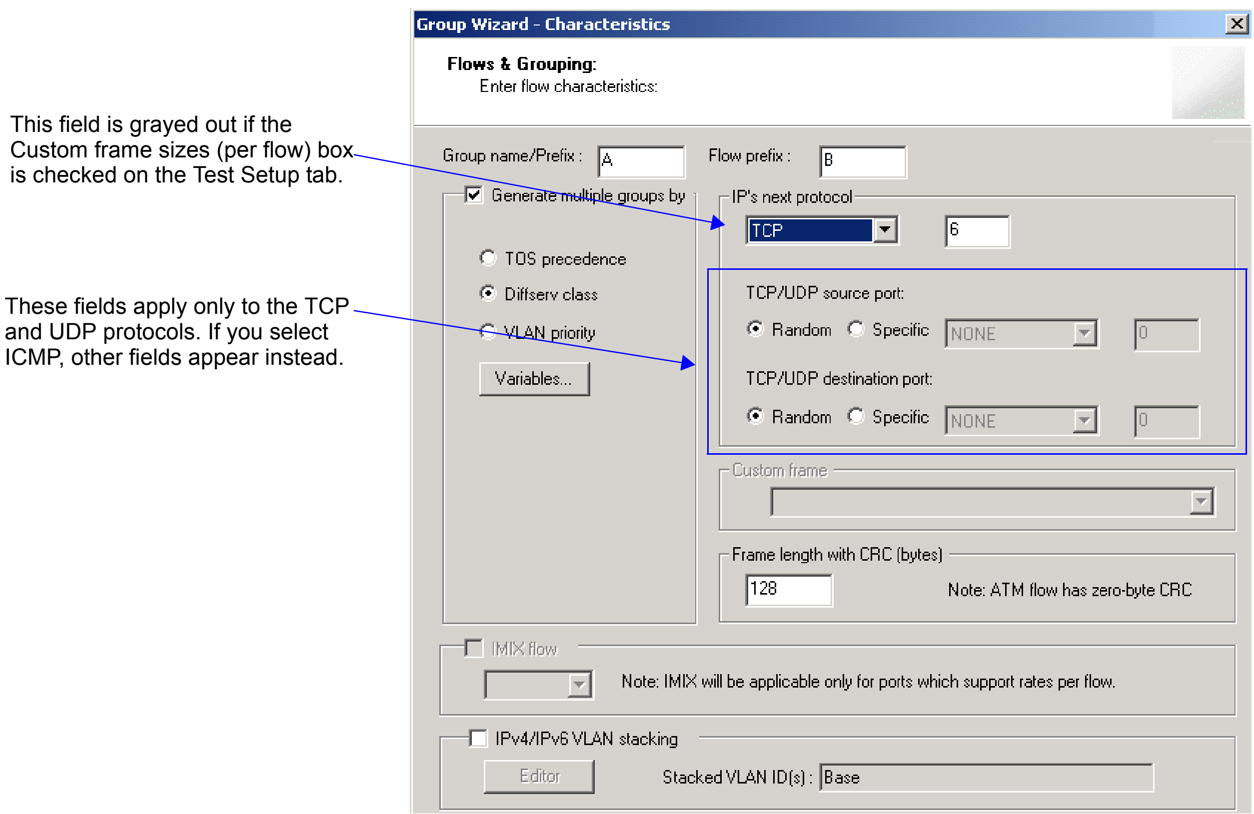

Figure 6-13. Group Wizard - Characteristics Page

Setting Rates per Flow by VLAN and IP Priority While Creating Flows

If you are creating flows and groups based on priority level, you can also set the rate of each flow created with the Group Wizard based on VLAN and IP priority.

As you create new flows using the Group Wizard, you can set the rate per flow based on VLAN and IP Priority, by using the *Rate Setting for ToS/Diffserv* dialog box (*Figure 6-14*). The rates per flow are set according to the flow's transmitting port/VLAN and IP priority (ToS or DiffServ). Once the flows are created, the rates are displayed on the *Groups* tab. From here, you can further modify the rates per flow individually or use the CoS Wizard to change the rates by VLAN and CoS.

All of the priority levels selected in the *Variables* dialog box are checked in the corresponding **Dx-V** (for Diffserv) or **Tx-V** (for ToS) check box in the *Rate Setting* dialog box. For example, if you selected Diffserv Precedences 0, 2, and 4, the following boxes would be checked: **D0-V**, **D2-V**, **D4-V**. The corresponding port column (D0-P, D2-P, and D4-P) would display 0 or the calculated percentage of the port rate.

You can copy down, fill increment, or fill decrement cells by highlighting and then right-clicking on the cells.

Notes:
- The sum of all Dx-P or Tx-P columns must not exceed 100% for a port. To verify that the port is not oversubscribed for all flows associated with the port, use the **Validate Rate** button on the *Groups* tab.
- The sum of all Dx-V or Tx-V cells in a row must not exceed 100% for a VLAN.
- If more than one flow uses the same port/subnet/VLAN/IP priority, the rate is evenly distributed among the flows.

Example: Assume the following:

VLAN Rate	D1-V	D1-P
100%	2%	2%
10%	2%	0.2%

If the test load is at 50%, the Dx-P (actual port load) is 2% of 50%.

Figure 6-14 shows an example of the *Rate Setting* dialog box. The D0-V is automatically checked since it was selected at the *Variables* dialog box.

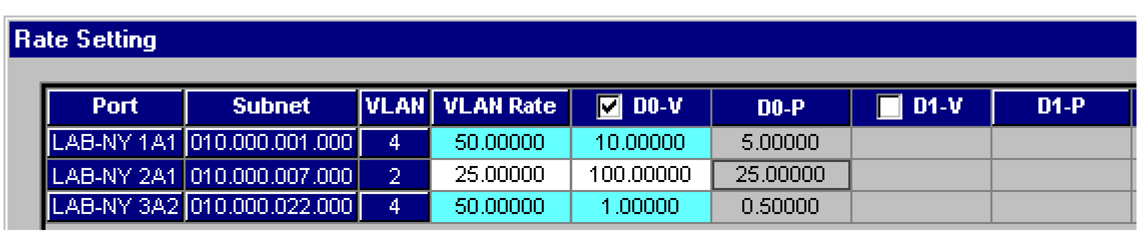
Rate Setting

Port	Subnet	VLAN	VLAN Rate	D0-V	D0-P	D1-V	D1-P
LAB-NY 1A1	010.000.001.000	4	50.00000	10.00000	5.00000		
LAB-NY 2A1	010.000.007.000	2	25.00000	100.00000	25.00000		
LAB-NY 3A2	010.000.022.000	4	50.00000	1.00000	0.50000		

Figure 6-14. Rate Setting Dialog Box Used to Set Rates per Flow Based on VLAN and IP Priority

Specifying Cyclic/Multiple Flows in the Wizard

Use the *Multiple Flows* (fourth) page of the Group Wizard to specify multiple SmartFlows at once and vary the fields in them, or specify a cyclic SmartFlow with variations all at once time. For any traffic pattern except waterfall, this page enables you to create either:

- Multiple non-cyclic SmartFlows with address variations (or a single, non-cyclic SmartFlow without address variation). This is the default. This option uses multiple streams if you vary the fields in the flows.
- A single, cyclic SmartFlow (with or without address variations). This option uses one stream, regardless of whether you vary fields. You can specify up to 65,000 variations of the source or destination IP address of one cyclic SmartFlow. You can also vary the MAC addresses and VLAN ID and VLAN Priority at the *Traffic* tab.

Generally, use cyclic flows if you want a large number of flows. Use non-cyclic flows for more combinations of address and/or port number variations.

One cyclic SmartFlow with no variations (*Number of Flows* field = 0) is the same as one SmartFlow with no variations.

You can use the *Traffic* tab to create a SmartFlow manually or to make a non-cyclic flow cyclic or vice-versa. To make *multiple* SmartFlows (cyclic or non-cyclic), use the right-click menu at the *SmartFlows* tab.

To vary more than one type of address (source and destination) at a time, you must manually specify the additional addresses for each flow at the *Traffic* tab. See *"Defining Individual Flow Traffic" on page 176* for more information.

Notes:

- The waterfall pattern of the Group Wizard does not utilize cyclic flows. (This page does not appear in the Group Wizard if you selected the waterfall traffic pattern.)
- SmartFlow does not support cyclic flows on the ML-5710 card.
- It is recommended that you use either all cyclic or all non-cyclic flows to more easily obtain a balanced traffic pattern.
- The contents of this page change if you chose to use custom frames on the *Group Wizard - Traffic Selection* page.

You specify incrementing fields for either cyclic SmartFlows or multiple, non-cyclic SmartFlows at the *Group Wizard - Multiple Flows* page. (See *Figure 6-15 on page 155*.)

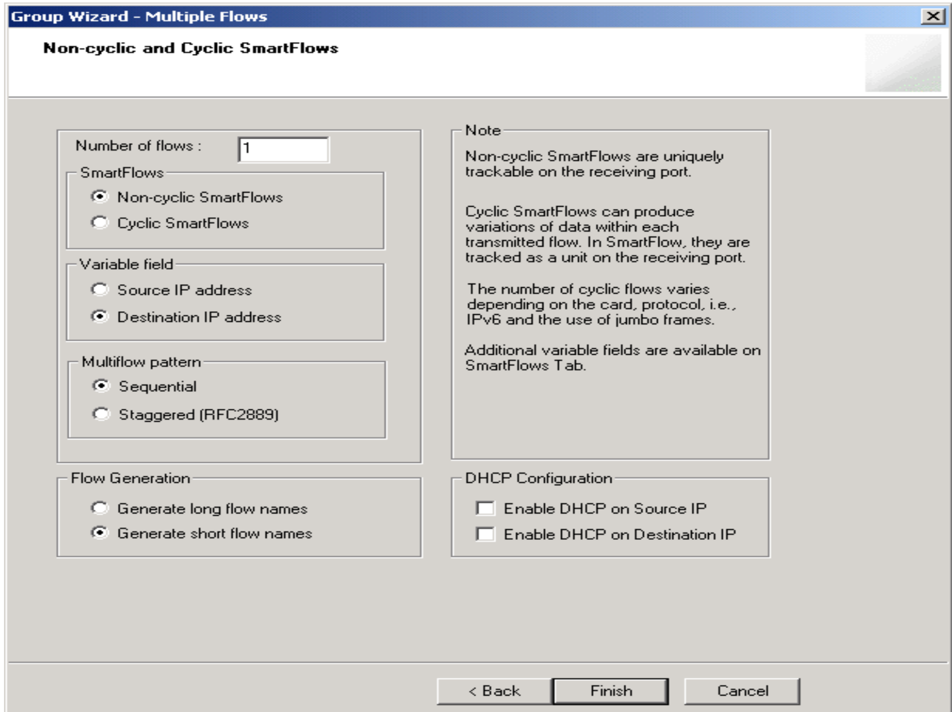

Figure 6-15. Group Wizard Multiple Flows Page

If you select *Non-cyclic SmartFlows* in the *SmartFlows* pane, the value in the *Number of Flows* field determines the number of flows created. Each flow generated will vary in whichever address that you specify; however, SmartFlow does not vary addresses *within* each SmartFlow, as it would with cyclic SmartFlows. To vary fields within flows, use the *Traffic* tab.

If you select *Cyclic SmartFlows* in the *SmartFlows* pane, SmartFlows will be cyclic. (This applies to both IPv4 and IPv6 flows). With cyclic SmartFlows, the specified source or destination address fields vary within each flow.

For field descriptions of the panes in this page of the Group Wizard as well as how to enable DHCP flows via the Group Wizard, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

When you click the *Finish* button, the flows generated from the Group Wizard are displayed on the *Groups* tab, as shown in *Figure 6-16 on page 156*.

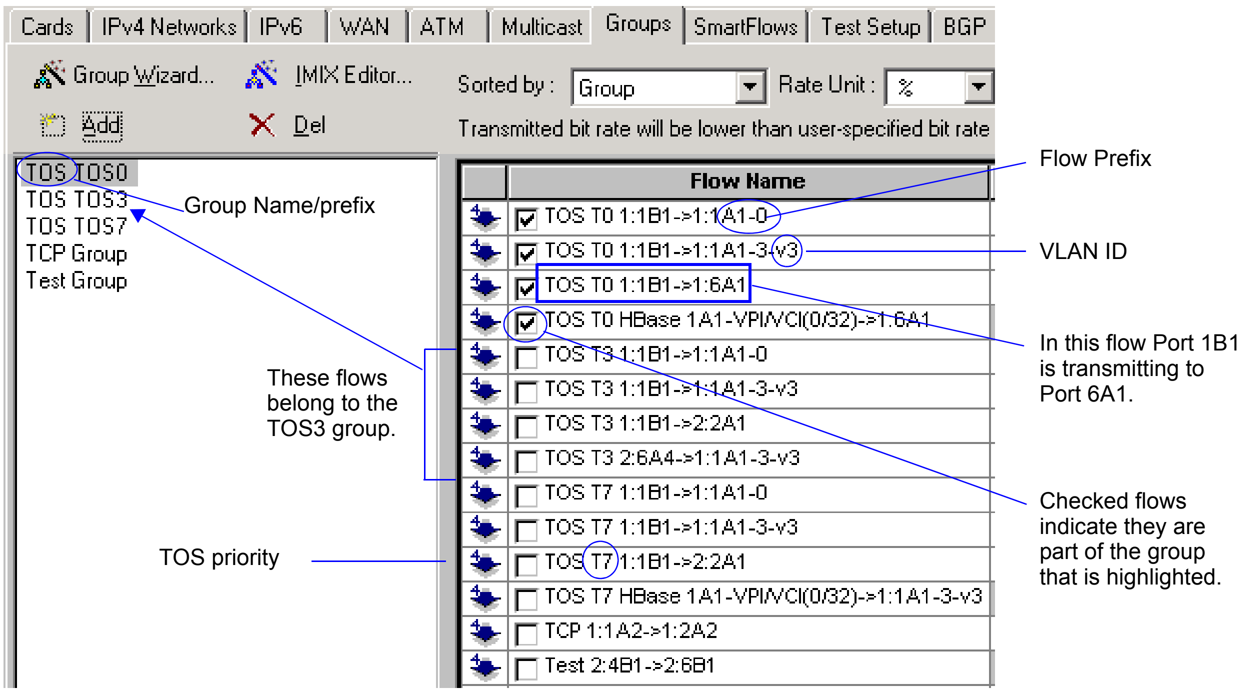

Figure 6-16. Example of Flows and Groups on the Groups Tab

Long and Short Flow Names

You can use the *Flow Generation* pane to select long or short flow names. By default, SmartFlow assigns a short name to each flow when it creates it. A short name simply gives the group name and port numbers to which the flow is transmitted and received. For example:

A 1->2 Group A, Port 1 to Port 2

In contrast, a long name for the same flow would appear as:

A Port 1->Port 2

Long names are useful when you have assigned aliases to ports—names that are meaningfully related to your test setup. For example:

A SF Lab 5A2->LA LAB 1A2

You assign port aliases on the *Traffic* tab of the *SmartFlows* tab. Double-click on a port-name entry and type in the desired name.

Specifying Waterfall Pattern Criteria in the Group Wizard

If you select the *Waterfall* traffic pattern, you can generate multiple groups in various ways, based on the parameters you select at the *Group Wizard - Characteristics* page. You can generate multiple groups of flows that are grouped by:

- Transmitting ports only
  Clear the **Priority** checkbox and select the **Transmitter** checkbox.
  See *"Generating Multiple Groups by Transmitter (No Priority)" on page 159*.
- Priority only
  Select the **Priority** checkbox and clear the **Transmitter** checkbox. (SmartFlow's underlying grouping method here is by transmitter, but it assigns priority to each transmitter group.) See *"Generating Multiple Waterfall Groups by Priority Only" on page 163*.
- Some flows and groups by priority only and some by transmitter only (without priority)
  Select both the **Priority** and **Transmitter** checkboxes.
  See *"Generating Multiple Groups by Transmitter and Priority" on page 166*.

Figure 6-17 shows an example of the waterfall pattern with ports selected in the *Group Wizard - Traffic Configuration* page.

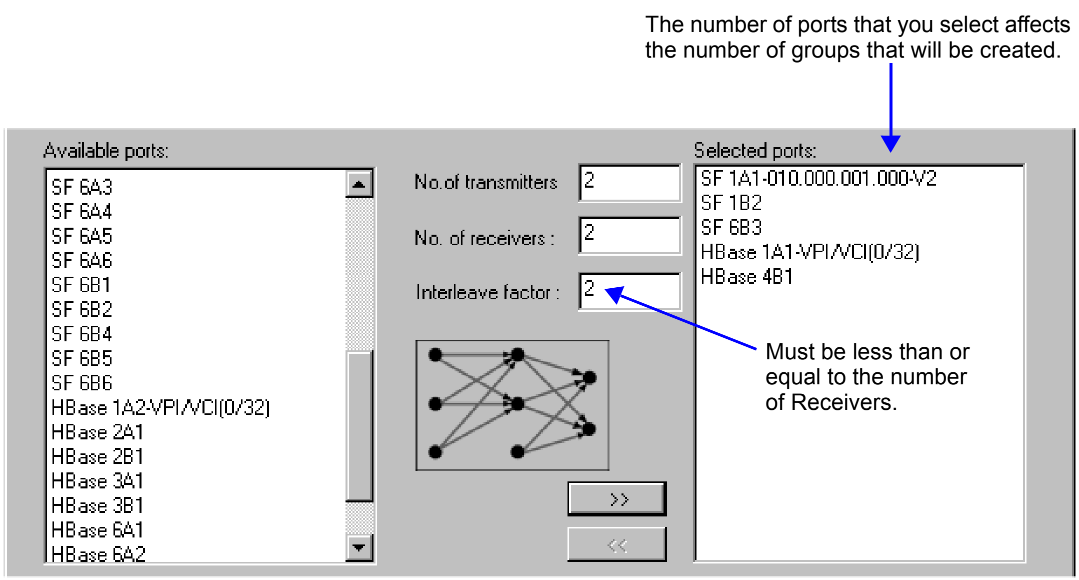

Figure 6-17. Example of the Waterfall Pattern with Ports Selected

The traffic pattern criteria selected in the example *Wizard* page shown in *Figure 6-17 on page 157* would result in the following flows:

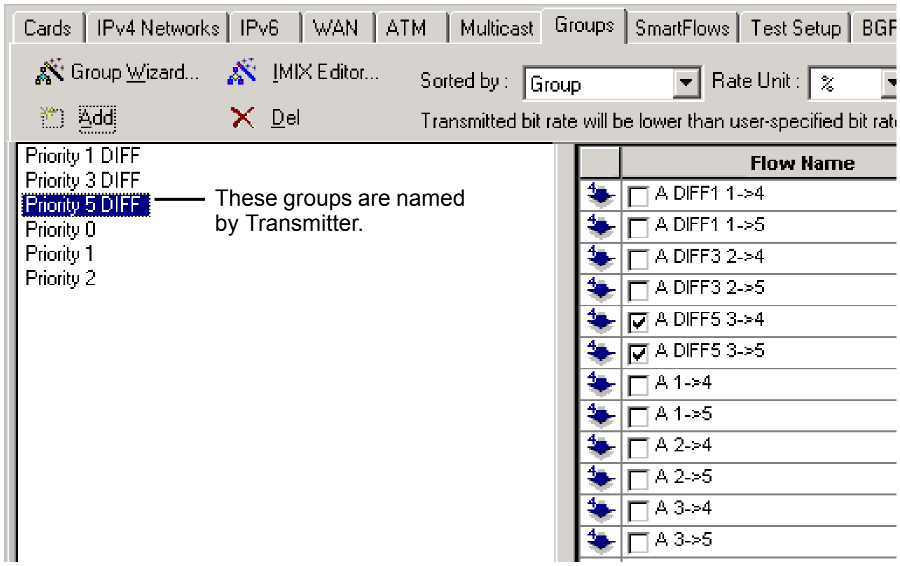

Generating Multiple Groups by Transmitter (No Priority)

Once you select the *Waterfall* pattern, you can elect to generate multiple groups based on the transmitters instead of the set of specified numbers of both transmitters and receivers.

Note: Grouping by transmitters does not increase the number of flows, only the way flows are grouped in test results.

You select this option on the *Group Wizard - Characteristics* page (*Figure 6-18*).

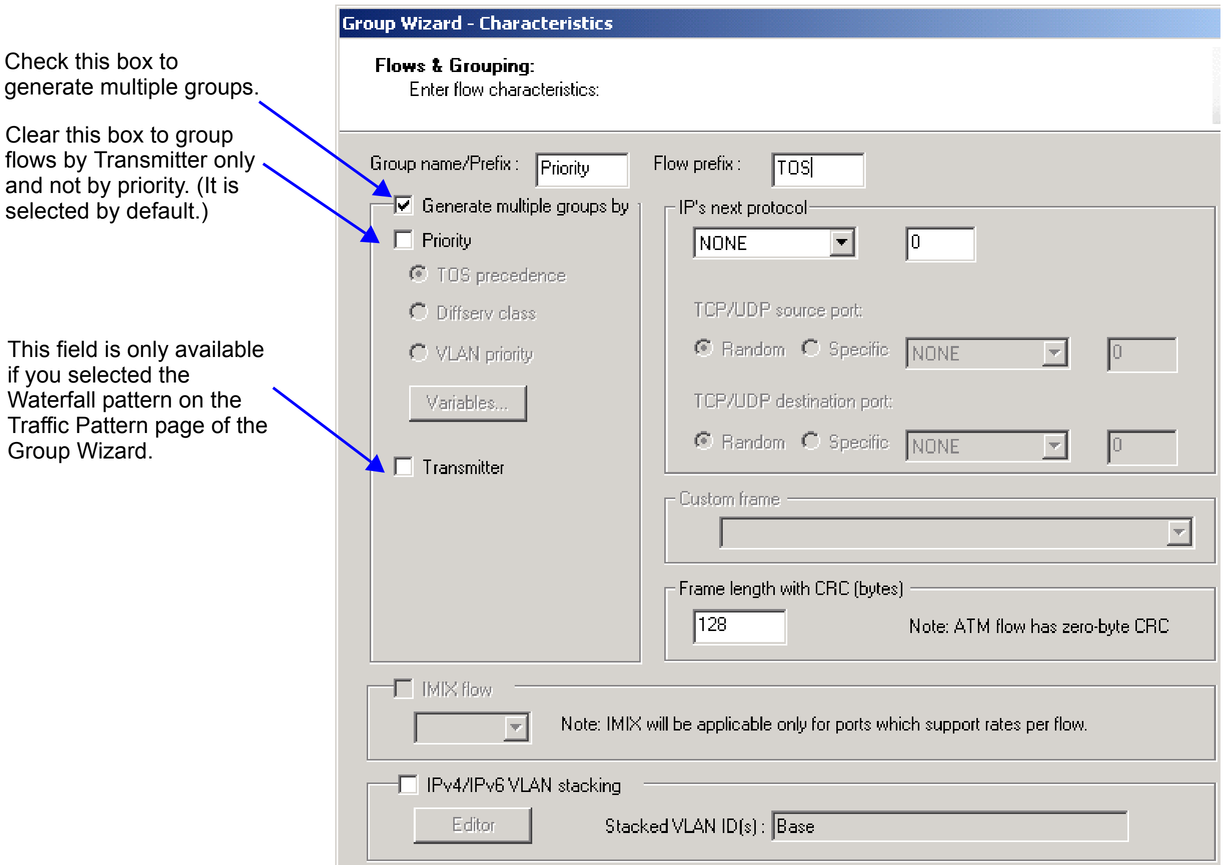

Figure 6-18. Characteristics Page of Group Wizard for Waterfall Pattern

The *Waterfall* pattern of transmitters and receivers remains the same as that illustrated in the previous examples, but the basis for creating the groups changes. The number of transmitter groups generated is equal to the number of transmitters that you specified for the *Waterfall* pattern criteria. Each transmitter group contains all of the flows generated by those transmitters.

This example includes *Figure 6-19* in which the groups were created without the *Generate multiple groups by* and *Transmitter* checkboxes selected. Once these checkboxes are selected, the groups created are shown here as thick ovals.

Ports Selected: **1- 8**
Number of Transmitters: **2**
Number of Receivers: **2**
Interleave factor: **2**
Generate multiple groups by: **selected**
Transmitter: **selected**

Because two transmitters were specified, SmartFlow creates two transmitter groups, each containing all the flows generated by each of the transmitters. For example, each transmitter whose flows are included in transmitter group 1 is the first transmitter in each (vertical) series of transmitting ports. Port 1 transmits, then port 3 transmits, and so on. Each of these ports is in the same group. The second transmitter in each series is port 2, then 4, and so on. These make up transmitter group 2.

In *Figure 6-19*, the circles (not ovals) show the groups that would be created if the *Generate multiple groups by* and *Transmitter* options were *not* selected.

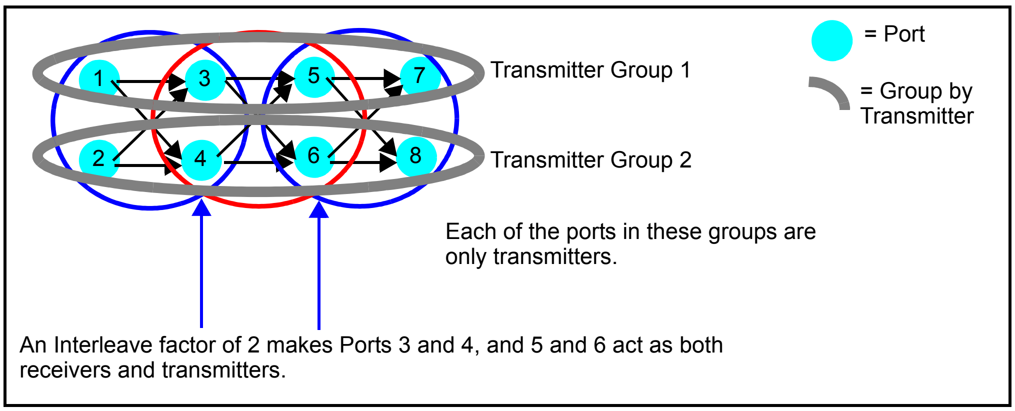

Figure 6-19. Waterfall Pattern with Groups Based on Transmitter

Based on the setup described in this section, the resulting flows created from the example configuration would look like this:

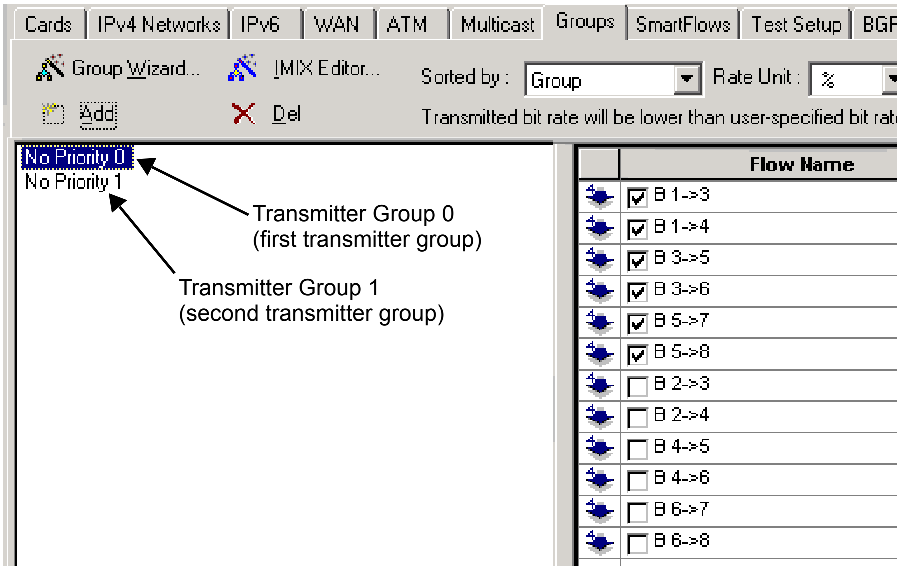

Note that the numbers in the group name are merely numbers of the groups, not priority levels. The first group is numbered 0.

If you selected criteria that resulted in a port only receiving and not transmitting, the port appears in flows and groups as a receiver only, and no transmitter group is based on it.

Figure 6-20 is another diagram that illustrates the concept of transmitter groups, with an interleave factor of 3.

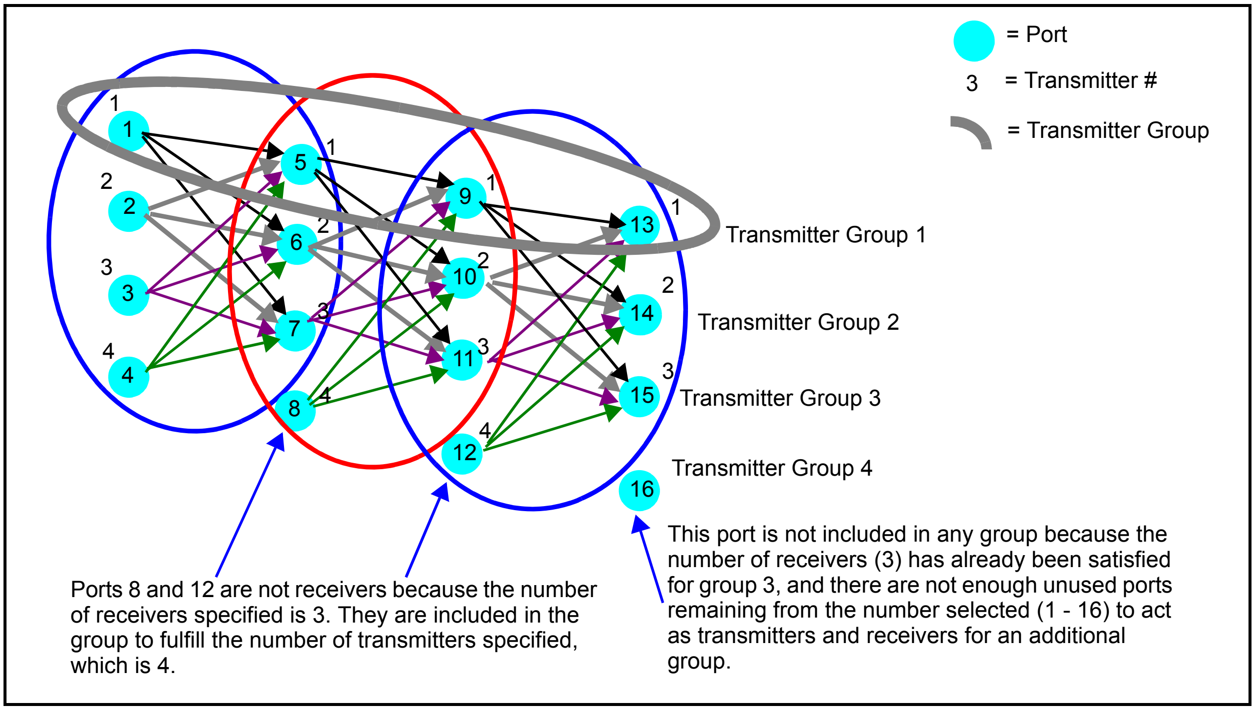

Figure 6-20. Waterfall Pattern with Groups Based on Transmitter and Interleave Factor of 3

For purposes of visual clarity, this diagram only encircles the first transmitter group with a thick line. However, the remainder of the transmitter groups are labeled as such.

Generating Multiple Waterfall Groups by Priority Only

Once you select the *Waterfall* pattern, you can elect to generate multiple groups based on priority in the pattern instead of on transmitters only or on the set of specified numbers of both transmitters and receivers in the pattern. You select this option on the *Group Wizard - Characteristics* page. (See *Figure 6-21*.)

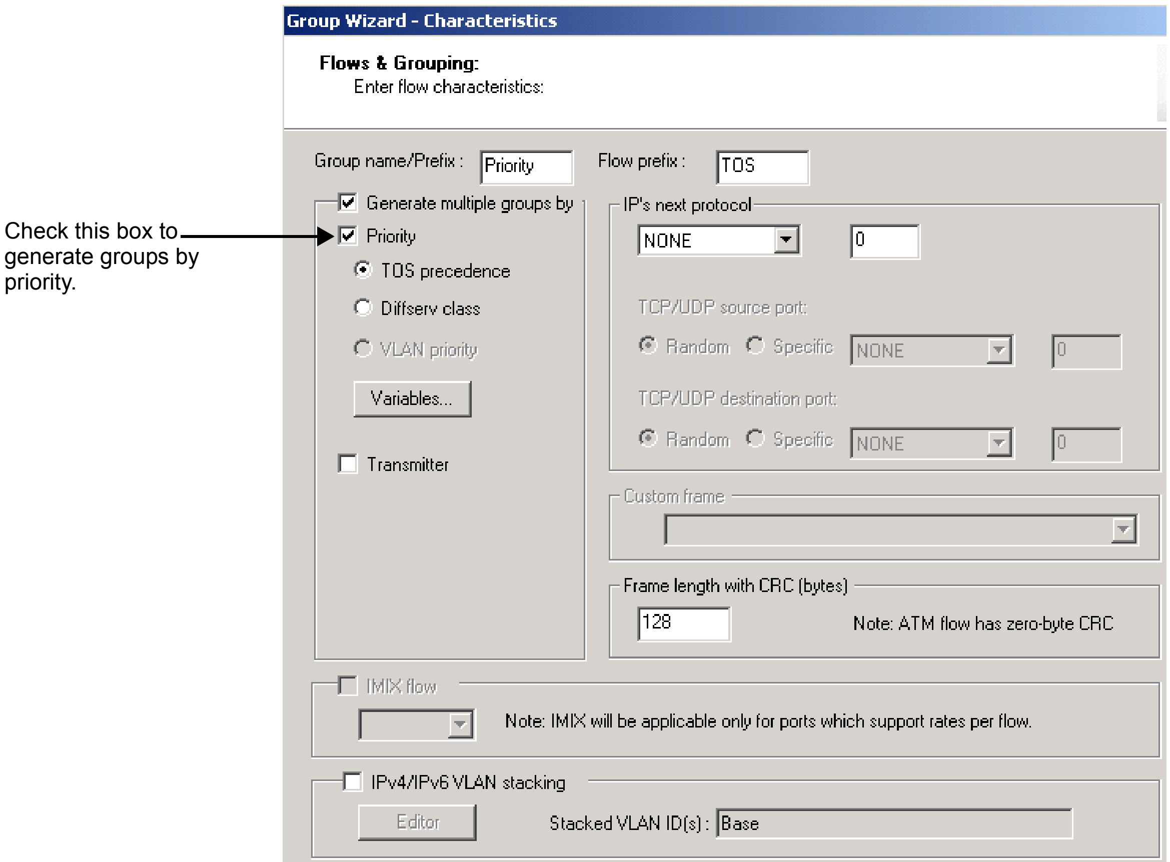

Figure 6-21. Group Wizard - Characteristics Page with Priority Selected for Waterfall Pattern

The *Waterfall* pattern of transmitters and receivers remains the same, except the basis for creating the groups changes.

When you decide to group by priority in the *Waterfall* pattern, SmartFlow maintains the concept of grouping by transmitter groups, even though the *Transmitter* checkbox is not selected. In addition, however, it assigns a user-defined priority to each transmitter group. Each transmitter generates flows of the same priority, and each priority group contains all of the flows generated by those transmitters.

Note: In the *Waterfall* pattern, grouping by a priority of 0 is the same as grouping by transmitter only (no priority).

Default Priority Groups

If you select the *Priority* checkbox but do not assign priority to each transmitter group (using the *Variables* button), SmartFlow assigns default priority levels to the groups according to the number of transmitters. For example, if you have three transmitters, the three resulting priority groups contain the following priority levels:

Group 1 = 0 priority level

Group 2 = 1 priority level

Group 3 = 2 priority level.

If the number of transmitters is nine or more, the ninth group's priority starts at 0 again.

Figure 6-22 illustrates the default priority levels that SmartFlow assigns to groups.

Ports Selected: **1- 10**
Number of Transmitters: **3**
Number of Receivers: **2**
Interleave factor: **2**
Generate multiple groups by: **selected**
Priority: **selected**
TOS Precedence: **selected**

Because three transmitters were specified, SmartFlow still groups all the flows generated by each transmitter into three groups, then assigns a default priority to each group, depending on the number of transmitter groups. The small numbers above each flow indicate priority. Notice how the priority groups correspond to transmitter groups even though the *Transmitter* checkbox was not selected.

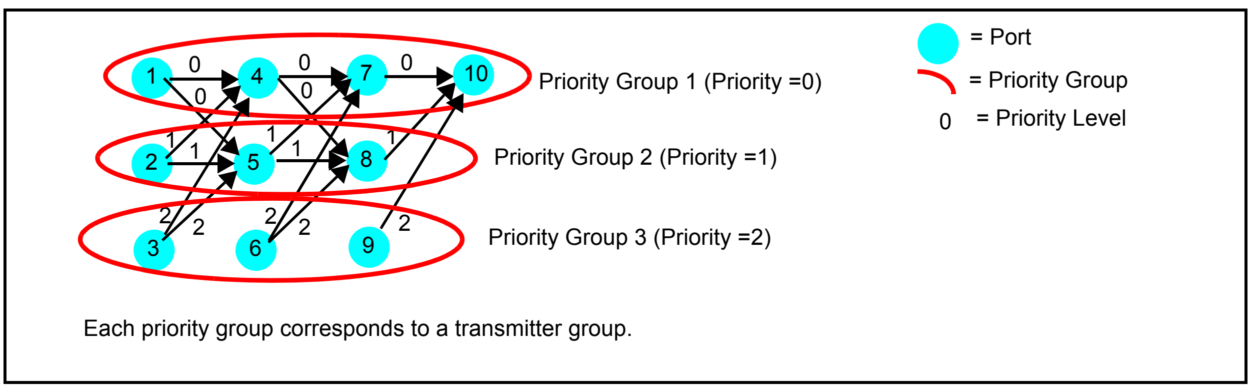

Figure 6-22. Waterfall Pattern with Groups Based on Default Priority Levels

User-defined Priority Groups

Figure 6-23 illustrates how SmartFlow generates multiple groups by user-defined priority levels.

Ports Selected: **1- 10**
Number of Transmitters: **3**
Number of Receivers: **2**
Interleave factor: **2**
Generate multiple groups by: **selected**
Priority: **selected**
TOS Precedence: **selected**

Let us say you assigned the following priorities to each transmitter group:

Transmitter Group 1: **Priority 0**
Transmitter 2: **Priority 3**
Transmitter 3: **Priority 0**

For information about how to assign priority to transmitters, see the online Help. Because three transmitters were specified, SmartFlow still groups all of the flows generated by each transmitter into three groups and then assigns a priority to each group. The small numbers above each flow indicate priority. Notice how the priority groups correspond to transmitter groups even though the *Transmitter* checkbox was not selected.

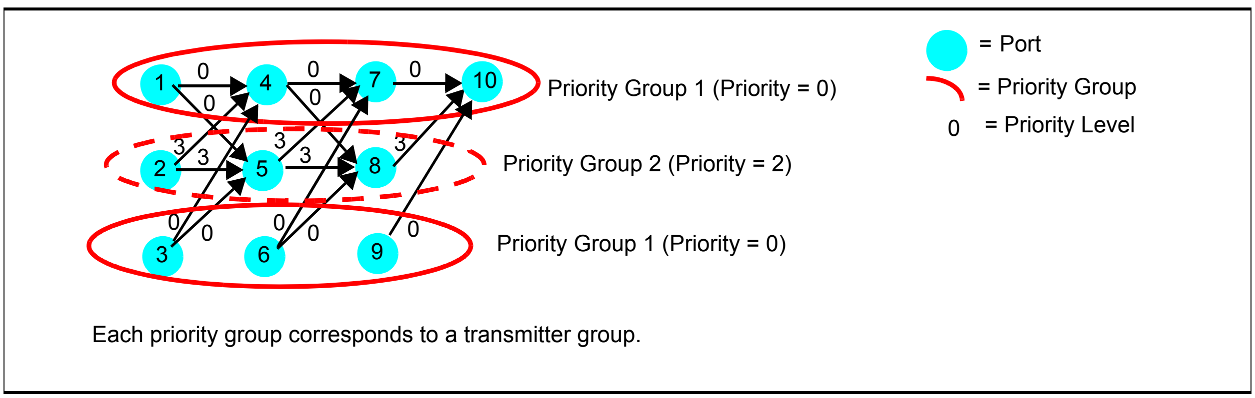

Figure 6-23. Waterfall Pattern with Groups Based on Priority

The resulting flows created would look like this:

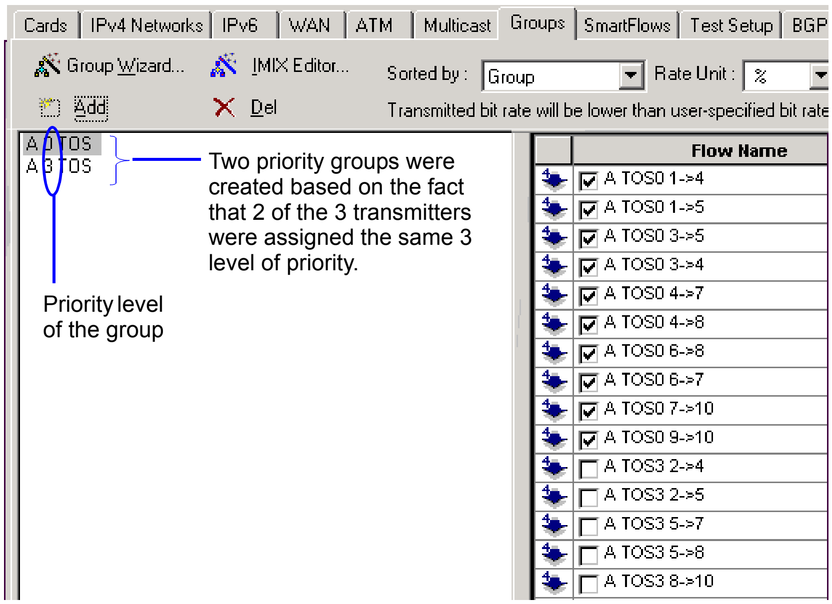

Generating Multiple Groups by Transmitter and Priority

Once you select the *Waterfall* pattern, you can elect to generate multiple groups based on both priority and transmitter in the pattern. SmartFlow generates some flows and groups based on priority only and an equal additional number based on transmitter only (without priority).

You select these options on the *Group Wizard - Characteristics* page. (See *Figure 6-21 on page 163*.)

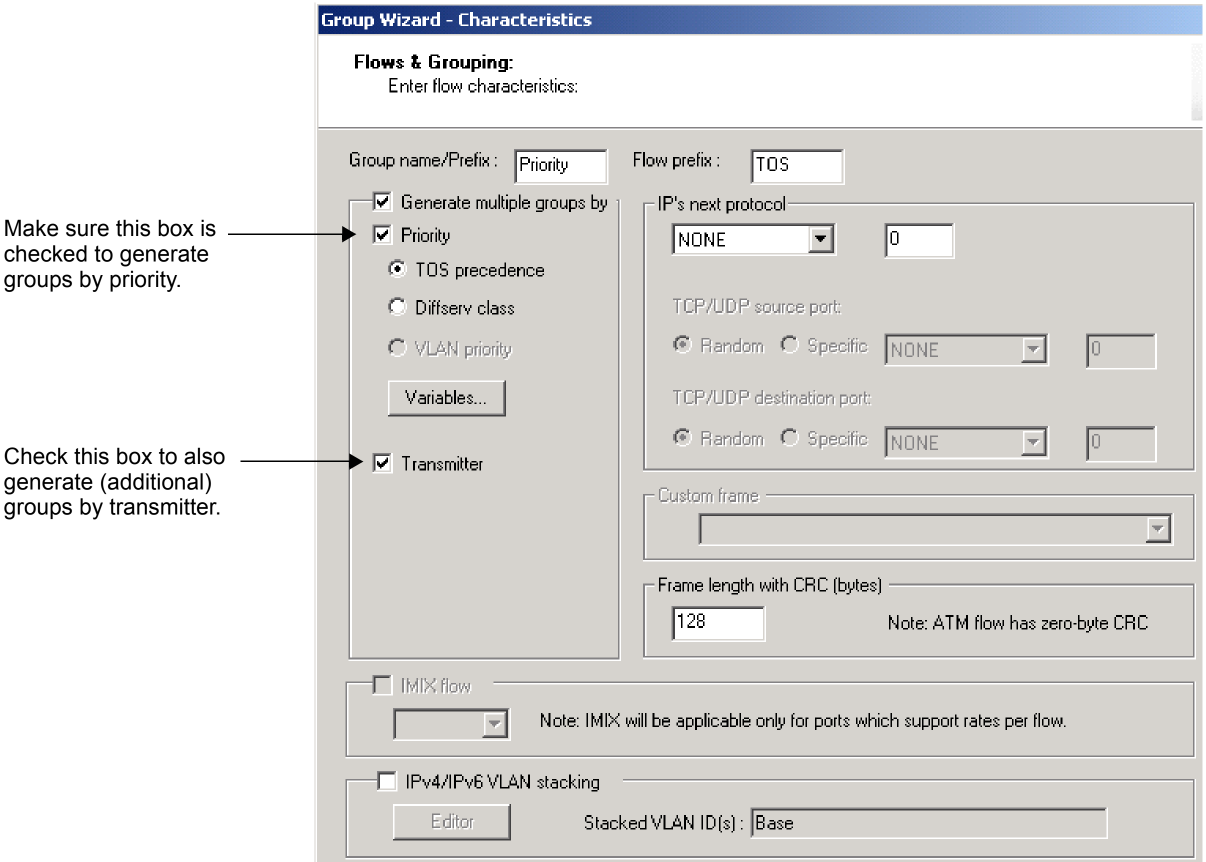

Figure 6-24. Group Wizard - Characteristics Page with Priority and Transmitter Options for Waterfall Pattern

The *Waterfall* pattern of transmitters and receivers remains the same, except for the basis for creating the groups changes.

Let us say you selected much of the same options as you did for generating multiple *Waterfall* pattern groups by priority only. (See *"Generating Multiple Waterfall Groups by Priority Only" on page 163*.)

Ports Selected: **1- 10**
Number of Transmitters: **3**
Number of Receivers: **2**
Interleave factor: **2**
Generate multiple groups by: **selected**
Transmitter: **selected**
Priority: **selected**

Let us say you assigned the following priorities to each transmitter group:

Transmitter Group 1: **Priority 0**
Transmitter 2: **Priority 3**
Transmitter 3: **Priority 7**

The resulting flows created would look like this.

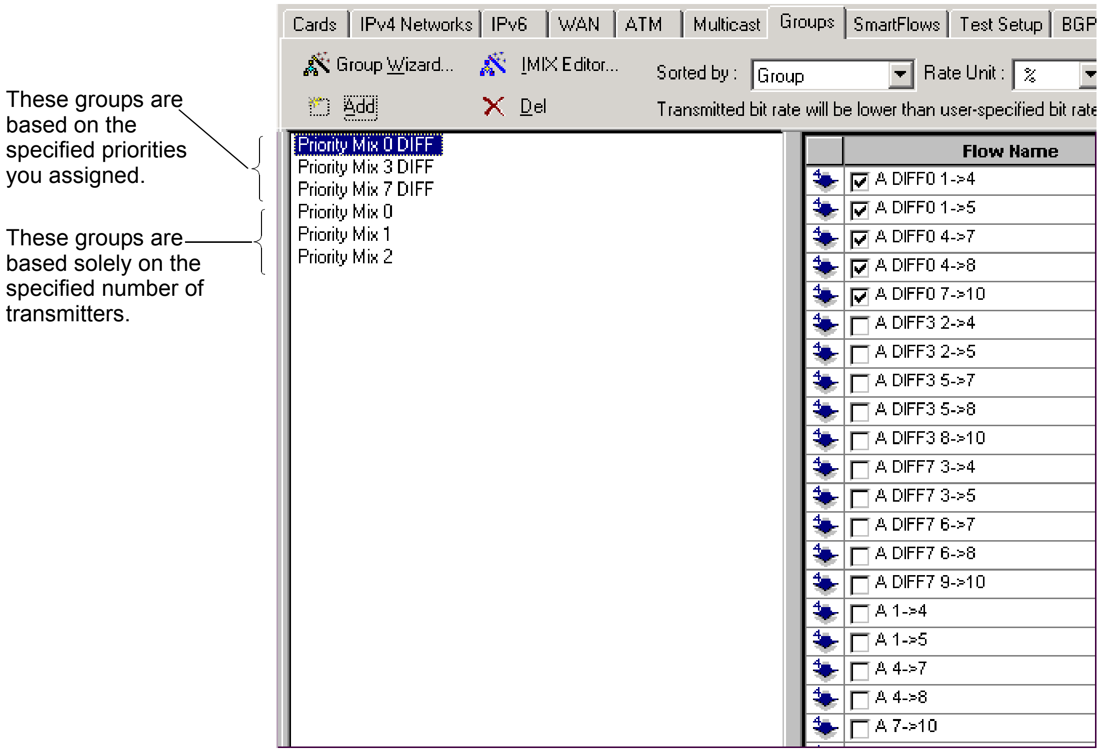

The flows and groups created are the same ones that would be created if you created them separately, by priority only and by transmitter only.

Modifying Flows and Groups Created with the Wizard

Once you use the Group Wizard to generate the flows and groups, you can view the attributes of the individual flows and modify them by clicking on the *SmartFlows* tab. For more information about modifying flows, see *"Copying, Modifying, and Deleting Flows and Groups" on page 205*.

This is best shown in the following examples. Let us say you use the Group Wizard to set up a group involving three ports in a mesh traffic pattern, as shown here in *Figure 6-25*.

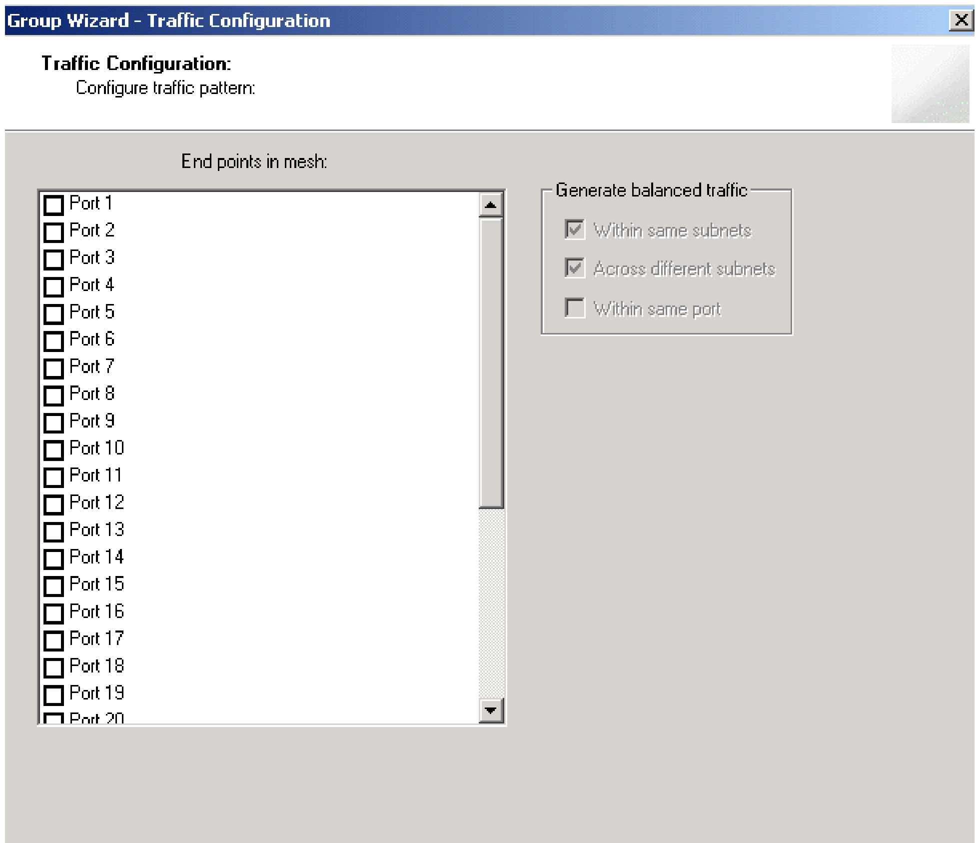

Figure 6-25. Fully Meshed Example in Group Wizard - Traffic Configuration

In this example, you named the group IP as they all contain IP flows. The Group Wizard generates the following flows.

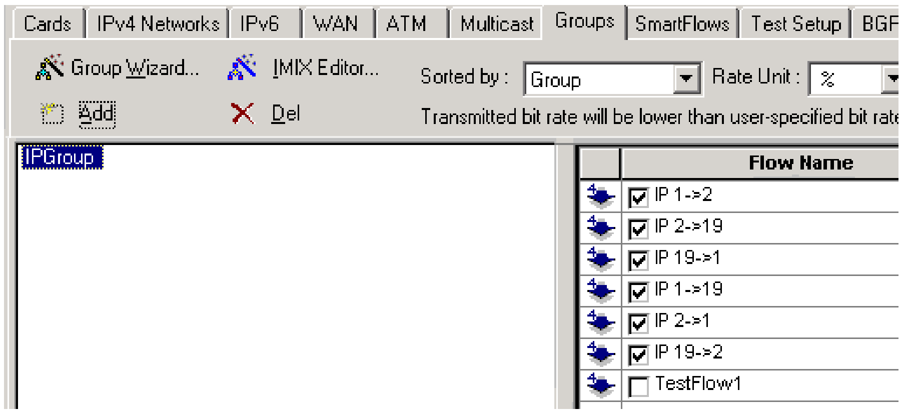

Notice that each flow in the group is automatically activated with a check. Now let us say you want another group that is similar to this group in that it uses *some* of the same ports and the same protocol (IP) and traffic pattern.

To create another group with similar flows:

1. Click the **Add** button on the **Groups** tab.
2. Enter a name for the new group.
3. Check the flows already generated by the Group Wizard that you want for the new group.

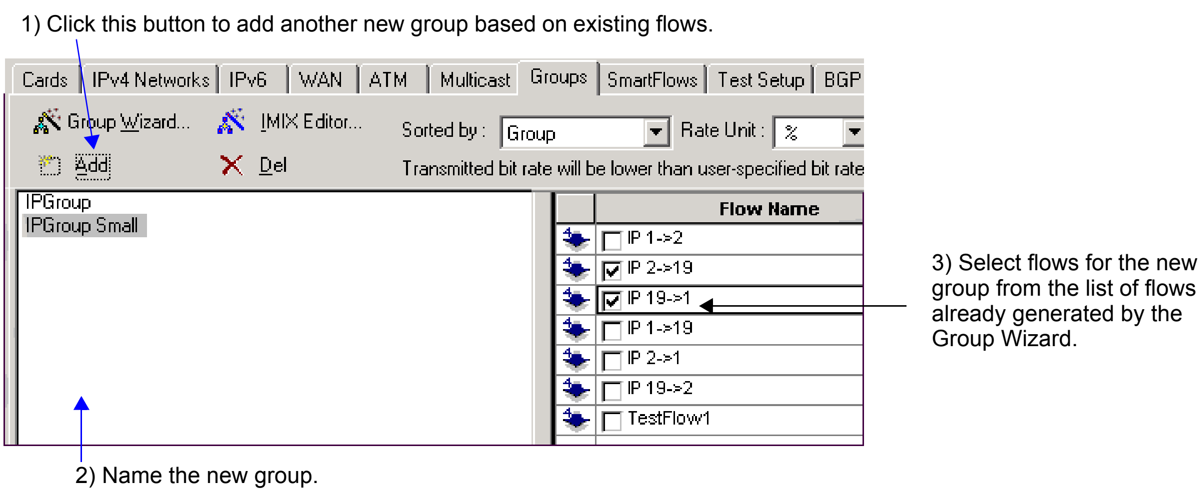

Let us say you want to further modify flow IP1->19 in this new group. To do this:

4. Click the **SmartFlows** tab.
5. Highlight the flow that you want to modify.

6 Click the **IP** tab. The IP parameters for that flow (such as *IP address*) are displayed.

Figure 6-26 shows the resulting flows. You can now modify these manually at the *SmartFlows* tab.

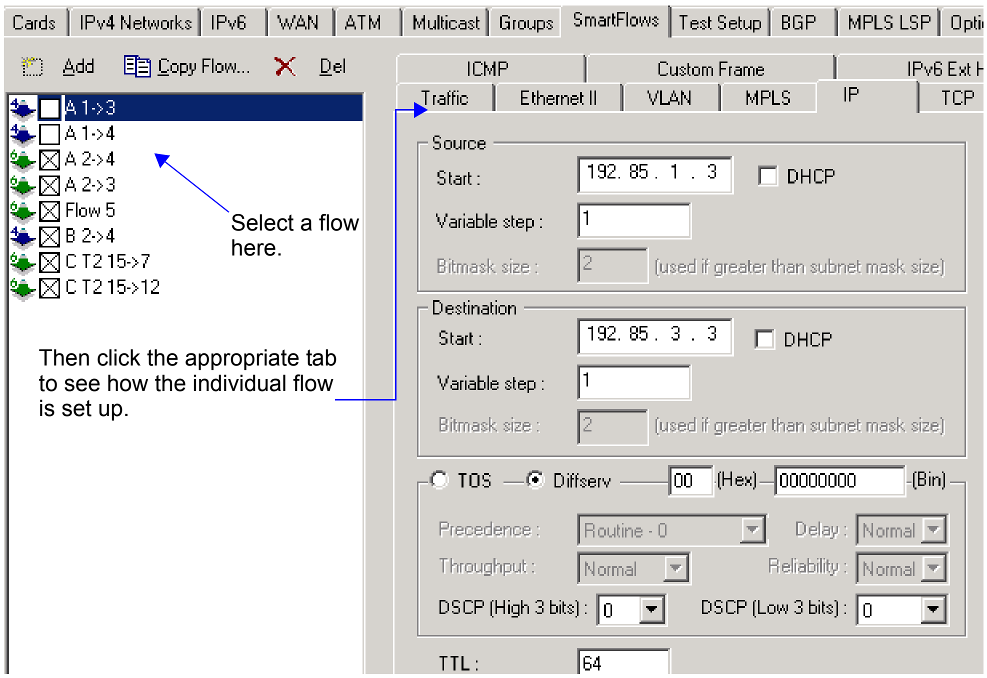

Figure 6-26. Manually Modifying a Flow at the SmartFlows Tab

SPIRENT
Communications

Chapter 7

Work with Individual Flows

This chapter explains how to modify flows and groups created with the Group Wizard as well as how to set up flows manually without using the Group Wizard.

In this chapter...

When to Manually Set up Flows

The Group Wizard should be your regular method of setting up groups and flows. (See *"About the Group Wizard" on page 132*.) But you can also set up or modify flows manually, one flow and group at a time, by using *SmartFlows* setup tabs.

Use the manual setup method when:

- You have already generated flows and groups using the Group Wizard but now wish to modify them.
- You want more control and flexibility when setting up flows and groups. For example, by setting up flows manually, you can vary the frame size within each flow.

Note: It is recommended that you connect to chassis before you set up the ports, networks, and flows. This ensures that flows and ports belong to the correct networks.

How to Manually Set up a Flow

If you do not use the Group Wizard to set up flows, you need to set up each flow of a group separately. For example, let us say you want to configure a one-to-many group (termed *backbone* in the Group Wizard). Port 1 transmits to Ports 5, 6, 7, 9 and 12. You have to create the five flows individually, then set up a group that includes these flows:

- `Port 1 to Port 5`
- `Port 1 to Port 6`
- `Port 1 to Port 7`
- `Port 1 to Port 9`
- `Port 1 to Port 12`

To set up a mesh test (all ports sending to all ports) using the same six ports, you have to set up 30 flows.

Note: If you have a large number of ports and/or you want to run a mesh test, it is best to use the Group Wizard.

To copy a flow, see *"Copying Flows" on page 205*.

Summary of Steps

Use these general steps to set up a flow manually:

1. In the main window, click the **SmartFlows** tab.
2. Click **Add** to add a new flow.
3. Enter a name for the new flow or use the default name, then press **Enter**.
4. On the *Traffic* tab, set up traffic information for the flow, such as the ports and protocol. For more realistic traffic, specify the fields to vary within the flow and the number of variations. See *"Defining Individual Flow Traffic" on page 176*.

   Your protocol selection determines which *SmartFlows* tabs you need to use next. For example, if you select *TCP,* you can specify fields on the *TCP* tab.

   When you wish to vary fields within a flow, you must use specific tabs to specify how the fields should vary.

   a. *If you want to vary the MAC addresses in the flow:*
   Click the **Ethernet II** tab. Specify the address(es). Define the amount by which to increment the address(es) and the byte offset to increment. See *"Varying MAC Addresses in a Flow" on page 181*.

   b. *If you configured any of the networks associated with the source port to send VLAN tagged frames*:
   Click the **VLAN** tab and select a priority level for the flow and additional VLAN tags for the flow. The ID is the VLAN ID number that you entered in the *Network* dialog box during port setup. See *"Assigning VLAN Priority for a Flow" on page 182*.

   c. *If you are testing MPLS:*
   Click the **MPLS** tab and set up label information.

   d. *If you selected IP as the protocol on the Traffic tab:*
   Click the **IP** tab to vary IP addresses, specify priority (if the device supports it), and set the number of hops (TTL) for the flow. See *"Defining IP Parameters for a Flow" on page 191*.

   e. *If you chose TCP or UDP as the protocol on the Traffic tab:*
   Click the **TCP** or **UDP** tab to specify the source and destination port numbers.

   f. *If you chose **SRC Port** and/or **DST Port** as a variable field(s):*
   Enter the number by which to vary the port type (in bits) in the appropriate *Variable step* field on the *TCP* or *UDP* tab, as appropriate. Incrementing the bits changes the type of port and thereby also the type of traffic. For more information about TCP/UDP, see *"Varying TCP/UDP Ports for a Flow" on page 202*.

   g. *If you chose ICMP as the protocol on the Traffic tab:*
   Click on the **ICMP** tab to specify the type and code for ICMP messages. For more information about ICMP, see *"Specifying ICMP Type for a Flow" on page 202*.

5. Repeat the appropriate steps starting with *Step 2* for each flow in the test.

Defining Individual Flow Traffic

Use the *Traffic* tab to define the basic properties of a flow, such as the source and destination ports and address values. To simulate more realistic network traffic, you can cause selected address and VLAN fields to vary with each transmitted frame. The number of variations is based on the value in the *Variable count* fields. You can simulate a complete network by varying the source and destination addresses with just one flow definition.

To open this tab, click the **SmartFlows** tab, then the **Traffic** tab (*Figure 7-1*).

For complete field descriptions, see the SmartFlow online Help. Press **F1** over the *Traffic* tab, or select **Help** from the menu bar. The online Help includes a procedure to set up traffic parameters for a flow.

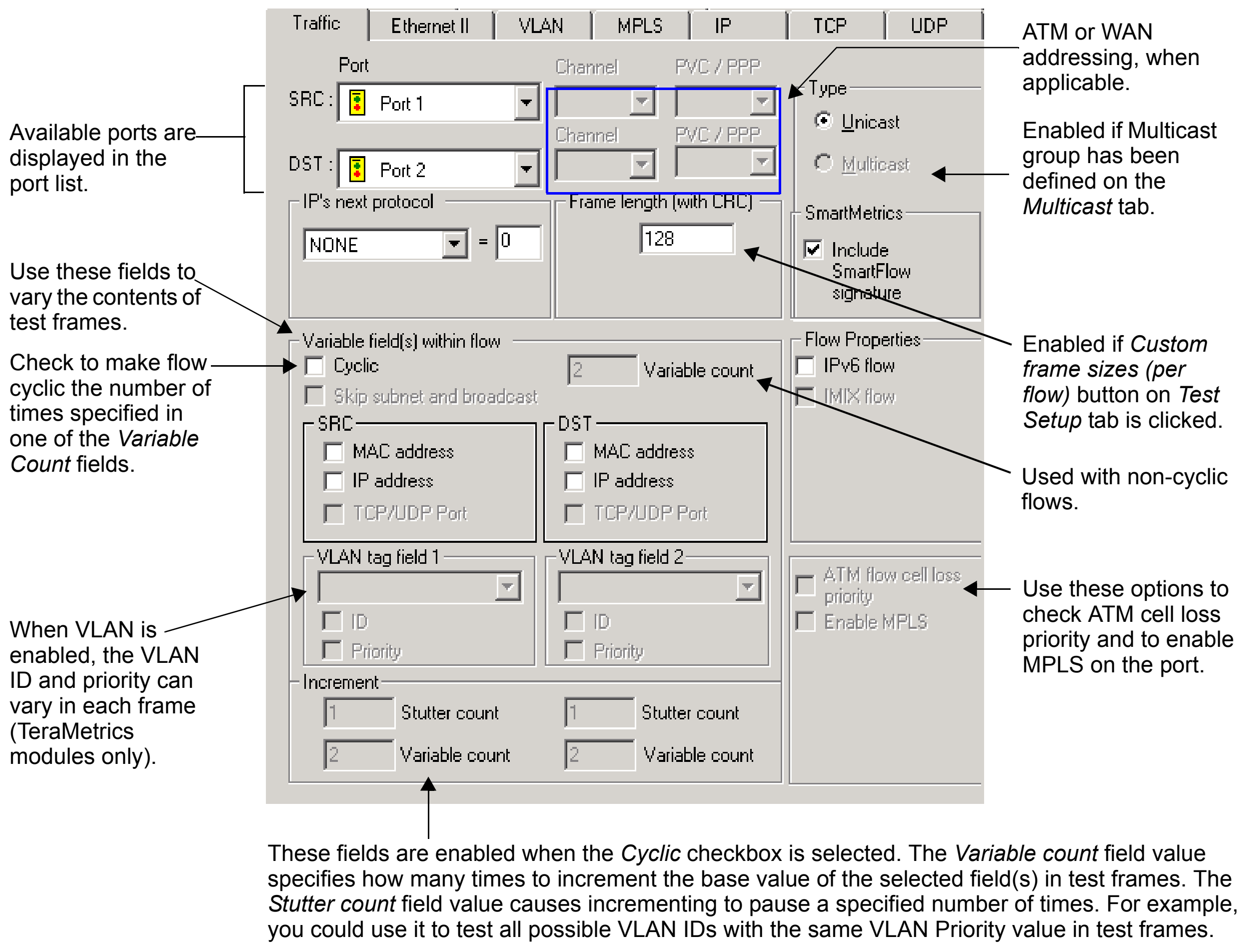

Figure 7-1. Traffic Tab

If you create a flow manually using the *Traffic* tab, the flow will always use the base network address of the port. To create flows that use a specific network address or VLAN, use the Group Wizard to create the flow or change the VLAN at the *VLAN* tab.

Notes:
- To enable the *Variable count* field, first check the field(s) you want to vary.
- Not all cards/modules support cyclic flows. See the online Help for complete information on cyclic flows and card/module support.
- The VLAN variable fields are active only when VLAN has been enabled for the port on the *IPv4 Networks* or *IPv6* tab and the VLAN ID has been set to a value other than **4096** (untagged) on the *SmartFlows>VLAN* tab. Refer to *Table 3-3 on page 62* and *Table 3-4 on page 63* for lists of these modules.

Important: If you create a multicast flow (for example, with the Group Wizard), use the *Traffic* tab to change the flow first to unicast, then back to multicast. The *DST* field does not retain the original IP address of the multicast group. By default, it is reset to the address of the first multicast group that you have set up. You must *reselect* the multicast group address if you do not want the default.

Varying Fields within a Flow

The options in the *Variable field(s) within a flow* pane enable you to vary the address values or VLAN tags within each flow as test traffic is sent. This creates more realistic test conditions. You can simulate a complete network by varying just the source and destination addresses with one flow definition.

For cyclic flows, the maximum number of fields that you can vary is two.

Refer to *Chapter 2, "Basic Theory,"* for a detailed description of SmartFlows and cyclic or non-cyclic flows. Briefly, *cyclic* flows are based on frame variations within one stream of traffic. They can give a much greater number of frame variations. To enable cyclic flows, select the **Cyclic** checkbox. *Non-cyclic* flows are based on multiple streams—one stream for each variation. They result in a comparatively smaller number of frame variations. To enable non-cyclic flows, clear the **Cyclic** checkbox. The card then generates one stream of traffic each time an address increments.

To enable the *Variable count* field, first check the field(s) that you want to vary.

Notes:
- For POS and WAN transmitting ports, the *SRC MAC address* and *DST MAC address* checkboxes are disabled. For ATM transmitting ports, they are disabled except when using bridged encapsulation.
- For cyclic flows, varying *SRC TCP/UDP Port* and *DST TCP/UDP Port* fields and VLAN fields are supported only on transmitting ports of TeraMetrics-based modules. See *"SmartBits Hardware Requirements" on page 59*.
- Cyclic flows sending ARP requests support varying *SRC IP*, *DST IP*, and *SRC MAC* fields.

Table 7-1 shows the *Variable field(s) within flow* pane combinations.

When you choose to vary fields in the flow, you must specify how to vary them. *Figure 7-2 on page 179* is a guide to which tabs to use to do this.

Table 7-1. Traffic Tab Variable Fields within Flow Pane Selection Options

If you select...	then...
IPv6 Flow	IPV6 traffic transmits from the selected port.
Cyclic	The *SRC TCP/UDP Port* and *DST TCP/UDP Port* checkboxes are disabled.
Variable count (*Cyclic* clear)	For each flow, *Variable count* streams are generated.
Variable count (*Cyclic* checked)	SmartFlow increments the selected field once each time the card transmits.
TCP or *UDP* as *IP's Next Protocol*	The *SRC TCP/UDP Port* and *DST TCP/UDP Port* checkboxes are enabled.
Multicast as the traffic type	The *DST IP address* and *DST MAC address* checkboxes are disabled.
An ATM Tx port	The *SRC IP address* and *DST IP address* checkboxes are enabled, along with the checkboxes that are enabled if TCP or UDP is the *IP's Next Protocol.* In addition, the *SRC MAC address* and *DST MAC address* checkboxes are enabled if bridged encapsulation is used.

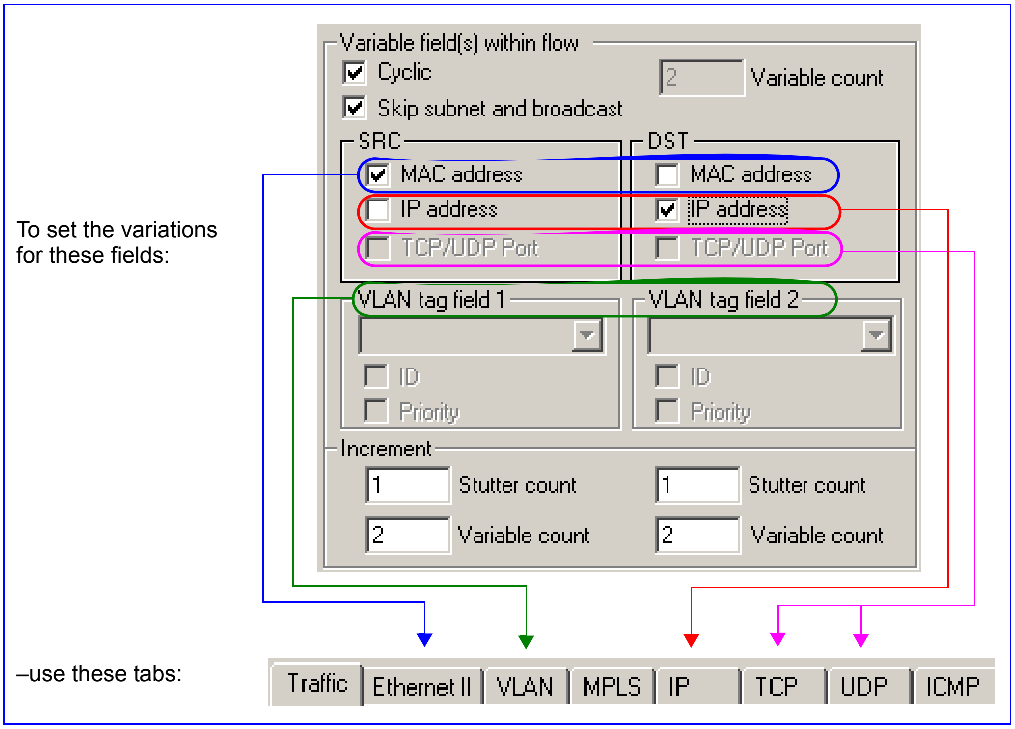

Figure 7-2. Defining Variations for Incrementing Fields

Changing Frame Sizes and Limiting Frame Length for Multiple Flows

If the Group Wizard was used to create flows each with the same frame size, you may now want to change the frame sizes and/or limit the frame length for random frame sizes. These processes are done using the *SmartFlows* tab as shown in *Figure 7-3*.

Important: Using frame size automation overwrites current frame sizes specified for flows. However, you can always manually change them back using this method.

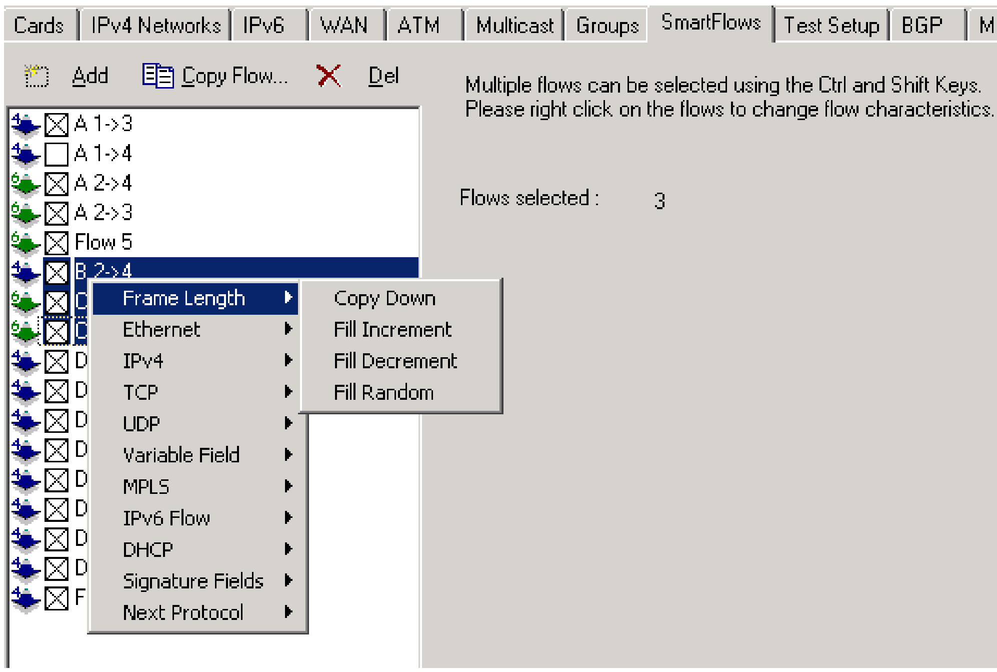

Figure 7-3. Changing Frame Sizes for Multiple Flows

To set the same frame size for multiple flows:

1. Click the **SmartFlows** tab, and select the **Traffic** tab.
2. In the flows list, highlight the first flow whose frame size is to be changed.
3. Enter the desired frame size in the **Frame length (with CRC)** field.
4. In the flows list, use the **Shift** or **Ctrl** keys to highlight the additional flows whose frame size is to be changed.
5. Right-click and select **Frame Length**>**Copy Down**.
   SmartFlow changes the frame sizes of all highlighted flows to the size specified for the first highlighted flow.

To enter random frame sizes varying between upper and lower frame length limits:

1. Click the **SmartFlows** tab.
2. In the flows list, highlight all flows whose random frame size is to be limited.
3. Right-click and select **Frame Length>Fill Random**. The **Random Frame Sizes (with CRC)** dialog box is displayed. To use legal sizes for a technology, use the default range (or a range that falls within that default range) for the type of card/module transmitting the flow as listed in the **Default size range** pane.

4 Enter a lower frame length limit in the **Minimum frame size** field, if necessary.

5 Enter an upper frame length limit in the **Maximum frame size** field, if necessary.
Important: Random frame lengths are generated within the capabilities of each individual flow.

6 Click **OK**.

Varying MAC Addresses in a Flow

If you selected *SRC MAC Address* or *DST MAC Address* as the variable field for a flow (on the *Traffic* tab), use the *Ethernet II* tab (*Figure 7-4*) to specify how the addresses should vary. You can also simply modify the source or destination MAC address once (instead of incrementing it) by changing the defaults on the *Ethernet II* tab. The MAC address in the *Start* fields does not need to correspond to the port's MAC address on the *Cards* tab.

To open the *Ethernet II* tab, click the **SmartFlows** tab, then the **Ethernet II** tab.

For complete field descriptions, refer to the SmartFlow online Help. Press **F1** over the *Ethernet II* tab, or select **Help** from the menu bar.

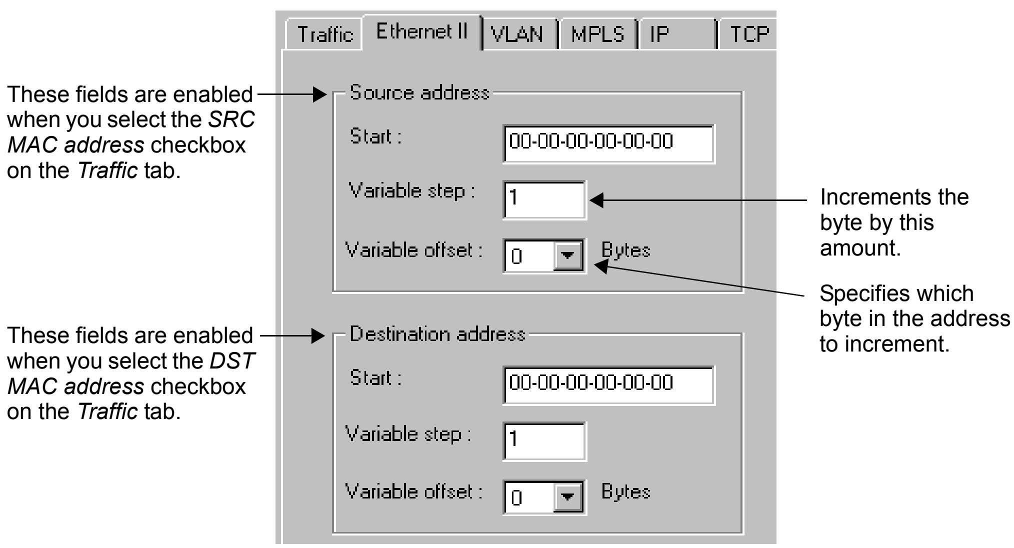

Figure 7-4. Ethernet II Tab

Note: All MAC addresses associated with a port must be unique. Make sure that the step and offset that you select do not result in duplicating an existing MAC address.

Examples

- If the MAC address is 00-00-02-00-00-02, the step is 1, and the offset is 0:

SmartFlow increments the MAC address to 00-00-02-00-00-03, 00-00-02-00-00-04, and so on, until the count specified on the *Traffic* tab is reached, after which it returns to 00-00-02-00-00-02.

- If the MAC address is 00-00-02-00-00-02, the step is 15, and the offset is 0:

  SmartFlow increments the MAC address to 00-00-02-00-00-11, 00-00-02-00-00-20, 00-00-02-00-00-2F, and so on, until it reaches the count specified on the *Traffic* tab, after which it returns to 00-00-02-00-00-02.

Assigning VLAN Priority for a Flow

You can use the *VLAN* tab (*Figure 7-5*) to modify the VLAN Priority and ID for a flow. This tab enables you to:

- View and set the VLAN ID and Priority for the flow's transmitting port.
- View the VLAN ID for the flow's receiving port.
- Add multiple VLAN tags (VLAN stacking) and edit VLAN tags.

To open the *VLAN* tab, click the **SmartFlows** tab, then the **VLAN** tab. All the fields on this tab are grayed, except the *ID* field, if VLAN is disabled for the flow's transmitting port or network, or if the *ID* field is set to **4096**. To enable VLAN for a port or network, use the *IPv4 Networks* tab or *IPv6* tab.

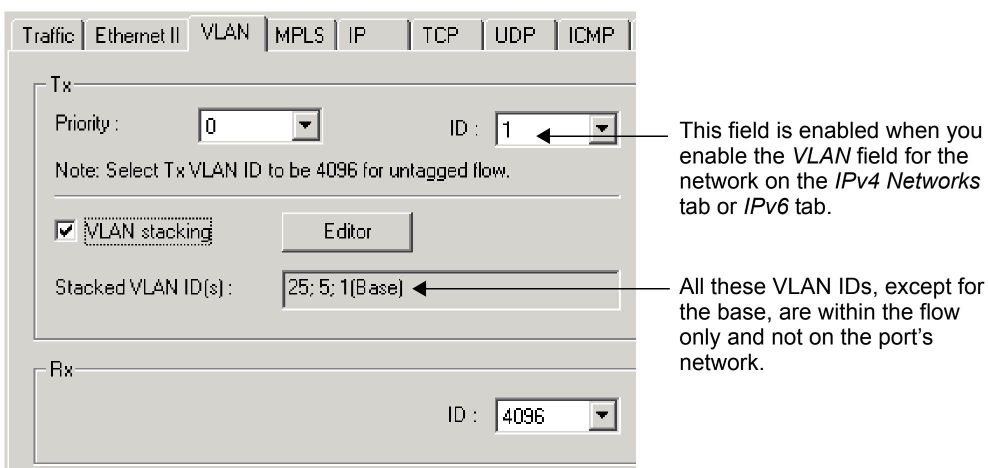

Figure 7-5. VLAN Tab

Setting VLAN Priority

In addition to TOS, Diffserv, or Traffic Class (IPv6) priorities, you can set VLAN priorities by using the *Priority* field. You can set different priorities within each VLAN tag as well as multiple priorities for the same VLAN from flow to flow. If the device does not support priority schemes, select **0** (default).

Once you set the VLAN priority of a flow, you can specify a rate for all flows in the same VLAN or with the same VLAN priority at the *Groups* tab.

Note: VLAN for a flow becomes disabled if you change an IPv6 flow with VLAN to an IPv4 flow and VLAN is not enabled for the transmitting IPv4 port. The same is true if you change an IPv4 flow with VLAN to an IPv6 flow, and VLAN is not enabled for the transmitting IPv6 port. If you enable VLAN after switching, the first VLAN ID on the port or network is used for the flow. VLAN also becomes disabled if you change the flow to have a VLAN ID of **4096** (untagged).

Selecting a VLAN ID

The *ID* field is the VLAN ID value assigned when setting up the port and network. The drop-down list shows only VLAN IDs that belong to the same network. To add more VLAN IDs to the list, use the *Network Wizard*, the *VLAN ID* Editor, or the *Multiple VLAN ID Wizard*.

The valid range for VLAN IDs is **1** to **4095**. A VLAN ID of **4096** indicates "untagged." The valid range for VLAN Priorities is **0** to **7**.

VLAN IDs of 4096

The *ID* field has a default value of **4096**. This means that the flow has an untagged VLAN: Either the card does not support VLANs, or VLAN is not enabled for the port or network. You can assign this value to a flow manually; even the transmitting port has VLAN enabled. This does not reset the *VLAN* field on the *IPv4 Networks* or *IPv6* tab.

Tip: By assigning **4096** for the *ID* field, a flow transmits untagged VLAN IDs—even if the flow's transmitting port is a tagged VLAN. This is because the value of 4096 is tied to the flow, not to the port.

Once you set the VLAN Priority of a flow, you can specify a rate for all flows in the same VLAN or with the same VLAN priority at the *Groups* tab.

Multiple VLAN Tags (VLAN Stacking)

If VLAN is enabled for the transmitting port of the flow, you can include multiple VLAN tags in the flow. Selecting the *VLAN stacking* checkbox opens the *VLAN Stack Editor*. You can use this to define a stack of VLANs tags for the flow.

The *Stacked VLAN IDs* field displays the VLAN tags created in the *VLAN Stack Editor* that will be included in the flow. Tags are displayed in the order in which they will appear in frames in the flow. This field is read-only. Base (for example, **5(Base)**) is the VLAN tag on the port or network. All tags are added to the left of the Base tag, so that it is the last in the stack.

Warning: If you disable VLAN stacking for a flow, assign an ID of 4096, or disable VLAN at the *IPv4 Networks* tab or *IPv6* tab, the existing definitions for the VLAN tag stack in the *VLAN Stack Editor* are deleted.

VLAN Stacking

VLAN stacking refers to the method of encapsulating one VLAN in another. VLAN stacking is also known as “Q-in-Q” or “super VLAN.” For an example of its application, see *“Why Use VLAN Stacking” on page 185*.

Use the VLAN Stack Editor (*Figure 7-6*) to create multiple VLAN tags to be included in flows. These tags are not added to the flow's transmitting port or network (on the *IPv4 Networks* tab or *IPv6* tab). The *Editor* enables you to define up to eight VLAN tags by defining the two bytes that make up the tag (Priority and VLAN ID) and the Ethernet Type field that precedes it.

The lower portion of the VLAN Stack Editor is a graphical representation of the order and location of where the tags will be added in a frame. Each tag shown represents the four bytes that make up a VLAN tag: E-type, Priority, and VLAN ID.

Note: VLAN tags added to a flow in a VLAN stack are not added to the network of the transmitting port.

For more information about VLAN stacking and how to use the VLAN Stack Editor, see the SmartFlow online Help. Press **F1** over the VLAN Stack Editor, or select **Help** from the menu bar.

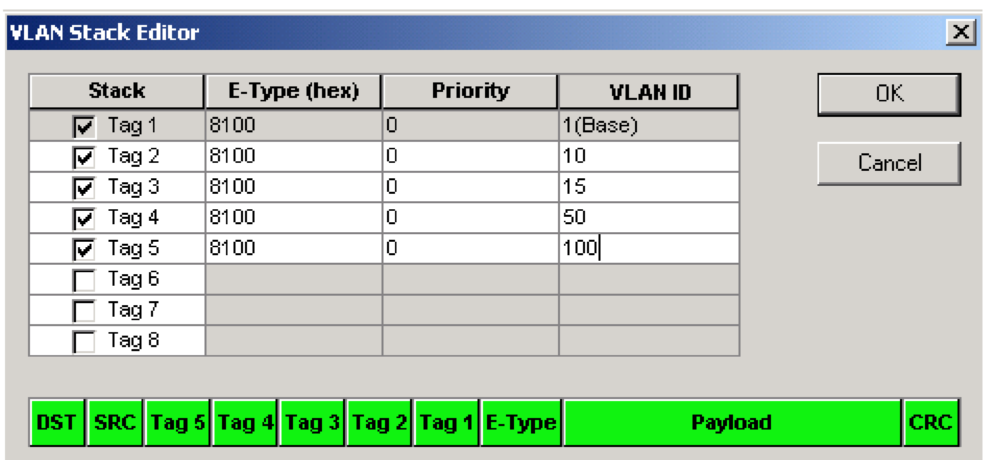

Figure 7-6. VLAN Stack Editor

Why Use VLAN Stacking

VLAN stacking enables a carrier (service provider) to partition a network among several national ISPs while allowing each ISP to utilize VLANs to their full extent. Without VLAN stacking, if one ISP provisioned an end user into VLAN 1 and another ISP provisioned one of their end users into VLAN 1, the two end users would be able to see each other on the network.

VLAN stacking solves this problem by embedding each instance of the 802.1Q VLAN protocol into a second tier of VLANs, i.e., the carrier assigns the first ISP "backbone VLAN". Within that backbone VLAN, the carrier uses a unique instance of 802.1Q VLAN tags to provide that ISP with up to 4096 VLAN identifiers. The second ISP is assigned a different backbone VLAN, within which another unique instance of 802.1Q VLAN tags is available.

Under 802.1q, a total of 4096 VLANs can be assigned. As described above, this number creates scalability problems in service providers and very large enterprise networks. With VLAN stacking, when VLANs are stacked (e.g., "two-deep" with two nested tags), carriers can assign 4096 VLANs within each VLAN. This allows a total of 16,777,216 (4096 x 4096) VLANs to be assigned.

Varying VLAN Parameters in a Flow

When testing with TeraMetrics-based LAN modules, you can vary the VLAN ID and Priority values in flows as test frames are sent. You do this by using the *VLAN tag field* options on the *Traffic* tab (*Figure 7-7*). You can do this for non-cyclic flows or cyclic flows. You can use cyclic flows with VLAN variations to transmit the entire range of VLAN IDs, regardless of the number of streams supported by the module. For example, LAN-33*xx*A XD modules can send up to 2047 stream definitions, less than the 4095 range of VLAN IDs, but up to 65,536 cyclic variations per stream, resources permitting.

This section explains:

- *"General Guidelines" on page 186*
- *"How to Use Varying VLAN Fields" on page 187*
- *"Varying VLAN Fields with VLAN Stacking" on page 190*.

For additional information on setting up VLAN parameters and varying VLAN fields, refer to the SmartFlow online Help. Press **F1** over the *Traffic* tab or *VLAN* tab, or select **Help** from the menu bar.

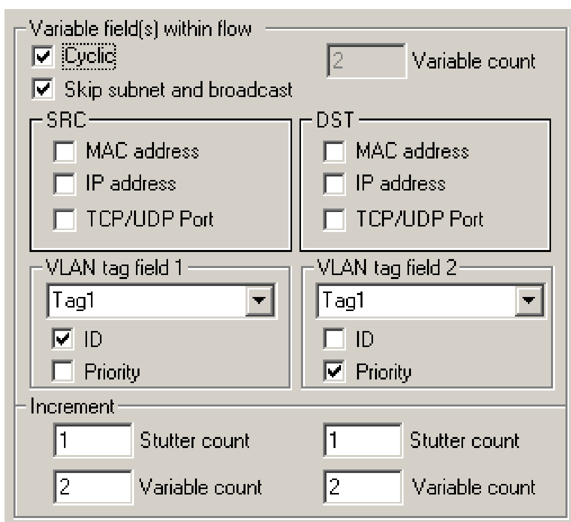

Use the VLAN tag fields to vary the VLAN ID and/or VLAN Priority in flows. When you enable Cyclic flows, each flow can increment through the entire range of VLAN IDs. (See *"General Guidelines"* for additional information.)

The two fields (tag field 1 and tag field 2) enable you to control the ID and Priority fields individually.

Figure 7-7. VLAN Options for Variable Field Definitions

General Guidelines

When using varying VLAN parameters, observe the following general guidelines:

- Varying VLAN fields are supported on TeraMetrics-based LAN modules.
- VLAN cyclic flows may be used with both IPv4 and IPv6 streams.
- VLAN IDs may be incremented through the entire range 1 to 4095. But if VLAN Priority is cycled on the same tag, the range for VLAN IDs is reduced to 255.

- Cyclic ARPing with VLAN variations is not supported. For example, if either the VLAN ID or VLAN Priority is varied, cyclic ARPing cannot be done for a selected IP address variable. ARPing is done only for the base stream.
- As with all cyclic variations, a *Variable count* field value of up to 65,535 is permitted.
- Only incrementing is allowed.
- The variable step is always **1**.
- Rollover of VLAN IDs is prevented. For example, if the base VLAN ID value is 4093 and the *Variable count* field value is 10, the resulting variations would range from 4093 to 4095, then restart at 4093. In effect, the count would be reduced from 10 to 3.

How to Use Varying VLAN Fields

The *VLAN tag field 1* and *VLAN tag field 2* options provide a way to set either the VLAN ID or Priority, or both. You can vary the VLAN fields alone or in combination with other variable fields.

The two fields can point to the ID and Priority values in one tag, as shown in *Figure 7-8*.

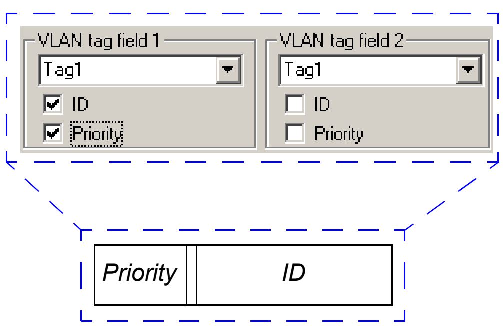

You can use the VLAN tag field 1 and VLAN tag field 2 options to set the ID and Priority value in one VLAN tag defined for the flow. With VLAN Stacking, the fields can also point to different tags. (See *Figure 7-13 on page 190*.)

Figure 7-8. VLAN Tag Fields Setting One VLAN Tag

In contrast, with VLAN stacking, they can point to the ID in one tag and the Priority in another, as shown in *Figure 7-9*. There are two fields to allow either the ID or Priority value (or both) to be selected as one of the two permitted variable fields in a cyclic flow.

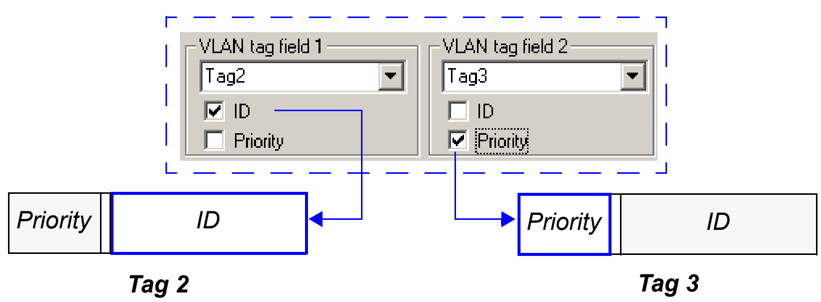

Figure 7-9. Varying Two Tag Fields with VLAN Stacking

Field Combinations with Cyclic and Non-cyclic Flows

When cyclic flows are enabled, you can set the ID or Priority in each VLAN tag field, but not both in the same field (*Figure 7-10*). When cyclic flows are disabled, you can set the ID and Priority values in any combination in either field.

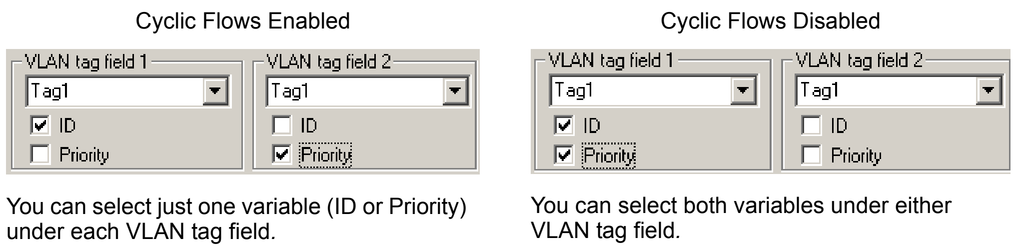

Figure 7-10. Selecting VLAN Tag Field Variables with Cyclic and Non-cyclic Flows

If VLAN Fields Are Not Active

The VLAN fields are not enabled if VLAN is disabled for the flow's transmitting port or network, or if the VLAN ID is set to **4096**.

To enable VLAN for a port or network, use the *IPv4 Networks* tab or *IPv6* tab.

The VLAN ID is the value assigned when you set up the port and network. The list displays only VLAN IDs that belong to the same network. To add VLAN IDs to the list:

- Use the *Network Wizard* on the *IPv4 Networks* or *IPv6* tabs.
- Use the *VLAN ID Editor*. Click the **SmartFlows>VLAN** tab, select the **VLAN stacking** checkbox, and click the **Editor** button.
- Use the *Multiple VLAN ID Editor*. When the *VLAN ID Editor* opens, click the **Add** button.

Number of Varying Fields

The allowed number of varying fields depends on whether flows are cyclic or non-cyclic:

- With cyclic flows, you can select one or two variable fields from among the options.
- With non-cyclic flows, you can select as many variable fields as desired.

Cyclic Flows Enabled   Cyclic Flows Disabled

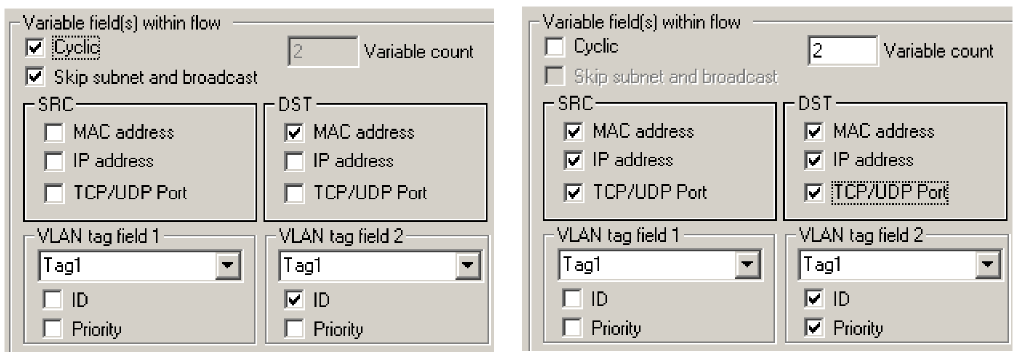

Cyclic flows permit two variable fields. For example, you could vary the VLAN ID or Priority in combination with an address field.

Non-cyclic flows permit any number and combination of variable fields.

Figure 7-11. Varying VLAN Tag Fields with Cyclic and Non-cyclic Flows

Varying VLAN Fields with VLAN Stacking

When you have set up VLAN stacking, the *VLAN tag field 1* and *VLAN tag field 2* drop-down lists enable you to vary the VLAN ID and Priority fields in different tags in the test frame. (See *"VLAN Stacking" on page 184*.)

Select a tag from the drop-down list, then specify which VLAN field should vary in that tag (*Figure 7-12*).

Figure 7-13 shows the VLAN ID and Priority fields selected in different tags in the VLAN stack.

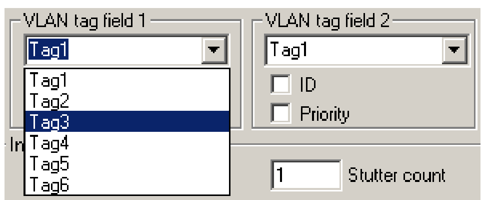

With VLAN stacking, select one of the configured tags from the drop-down list. Then select the variable field(s): ID and/or Priority.

Figure 7-12. Varying a Selected Tag with VLAN Stacking

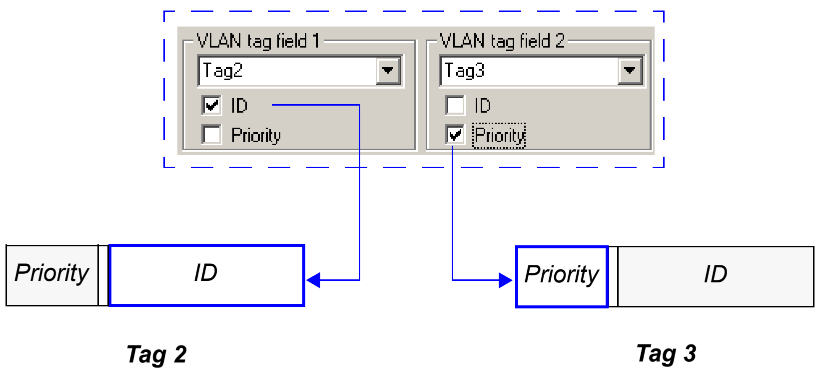

Figure 7-13. Varying Two Tag Fields with VLAN Stacking

Defining IP Parameters for a Flow

For testing based on the *TOS*, *Diffserv,* or *Traffic Class* field, use the *IP* tab to specify priority scheme information. You can also use the *IP* tab to vary the source and/or destination IP addresses in the flow. You can modify the default address or specify how to increment the address. Creating many variations simulates realistic network traffic. You set the variable step here. The number of variations is based on the value in the *Variable count* fields on the *Traffic* tab.

To open the *IP* tab, click the **SmartFlows** tab at the main window, then the **IP** tab. The tab shows different fields for IPv4 and IPv6 flows.

- *Figure 7-14* shows the tab fields for IPv4 flows.
- *Figure 7-15 on page 192* shows the tab fields for IPv6 flows.

For field descriptions and a procedure to define IPv4 or IPv6 flows, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

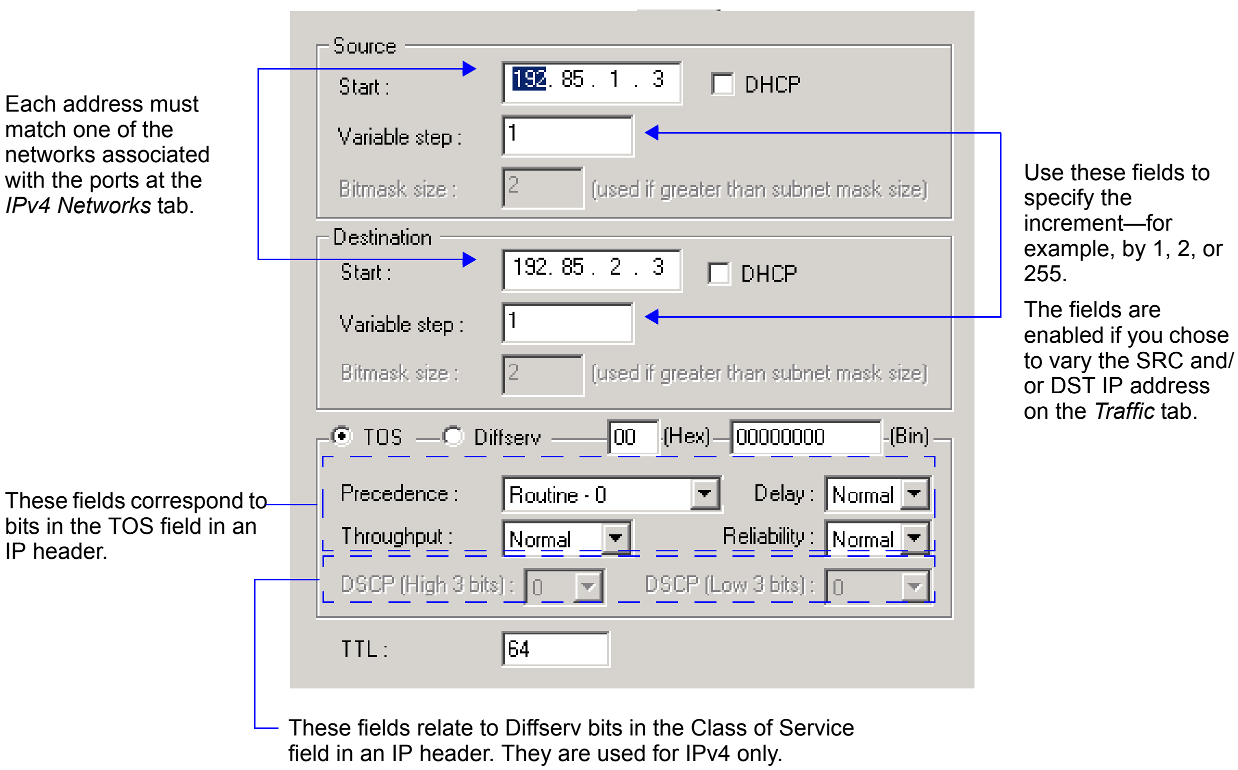

Figure 7-14. IP Tab for IPv4 Flows

If the flow is an IPv6 flow, the *IP* tab displays different fields as shown in *Figure 7-15*. Your test can send IPv6 flows when:

- The source SmartBits port is IPv6-capable.
- IPv6 is enabled for the port (*IPv6 Capable* is selected on the *Cards* tab).
- The *IPv6 flow* checkbox on the *SmartFlows>Traffic* tab is selected.

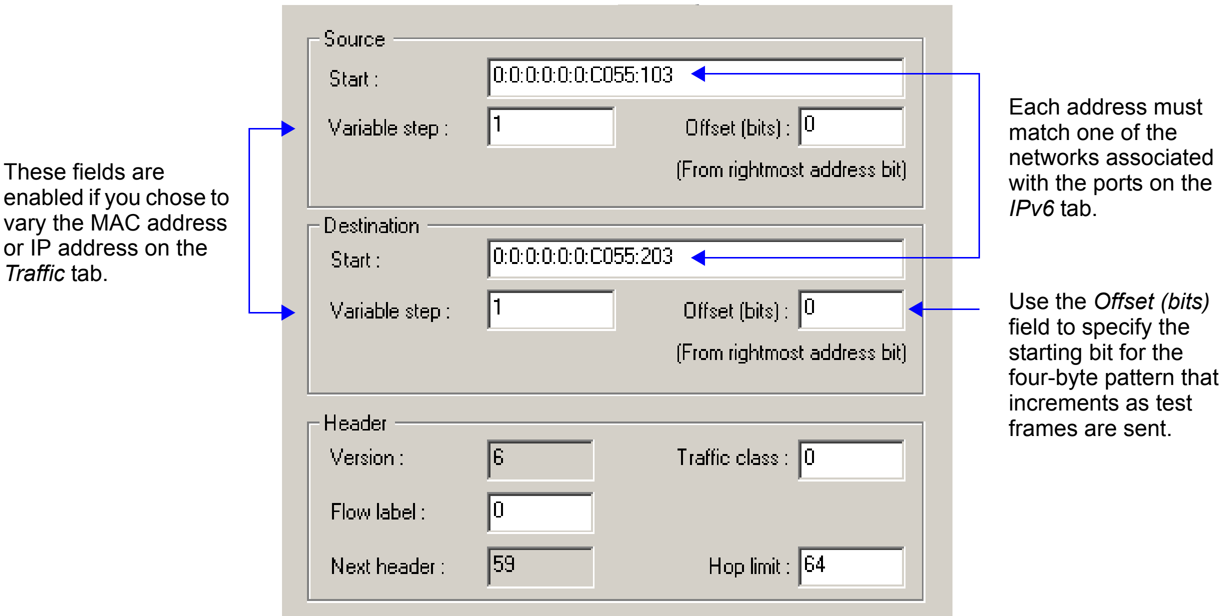

Figure 7-15. IP Tab for IPv6 Flows

Varying IP Addresses

To make traffic more realistic, you can vary the IP addresses of flows.

Changing IPv4 Addresses for Multiple Flows

If you used the Group Wizard to create multiple IPv4 flows, you may now want to change the IP addresses manually. Refer to the online Help for information about how to do this.

Figure 7-16 illustrates one of the dynamic menus used to change addresses. The addresses increment/decrement or randomly fill according to the bit mask that you select. (See *Table 7-2*.)

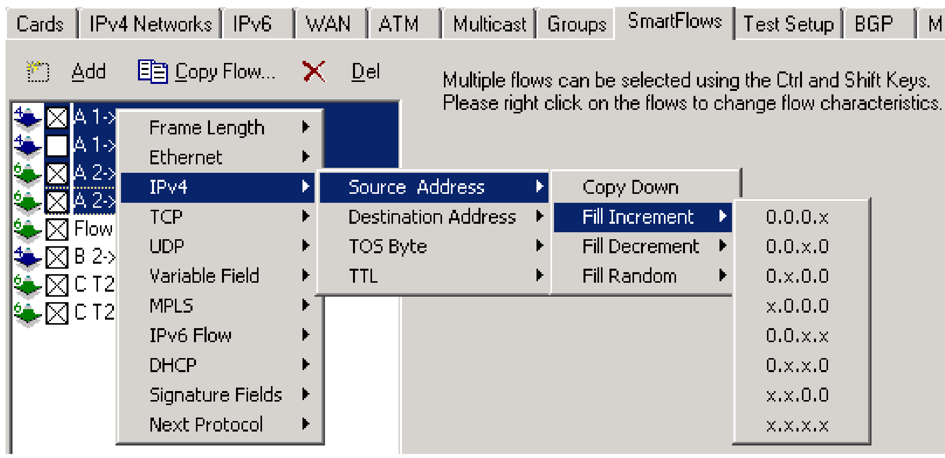

Figure 7-16. Changing IPv4 Addresses for Multiple Flows

Table 7-2. How IPv4 Addresses are Incremented/Decremented or Randomly Filled

Bitmask Selected	Result
Single x byte (0.0.0.x, 0.0.x.0, x.0.0.0)	SmartFlow changes only the byte position represented by x. For example, 0.0.0.x changes the right-most byte.
Double x.x bytes (0.0.x.x, x.x.0.0, 0.x.x.0)	SmartFlow first increments or decrements the less significant (right-most) byte position represented by the x pair up to 255 or down to 0, then changes the more significant byte position.
All bytes x.x.x.x	SmartFlow changes all four bytes simultaneously.

Note: With IPv4 flows, if you select a bit mask that affects the network portion of the address, SmartFlow sends an error message.

How IPv4 Flow IP Addresses Relate to the IPv4 Network Setup

The network portion of a flow's IPv4 address (source or destination) must belong to one of the same networks that is associated with the port on the *IPv4 Networks* tab, once you apply the subnet mask.

Example: Let us say a port has four networks associated with it:

192.085.001.000
010.000.001.000
010.000.002.000
010.000.003.000

If the subnet mask is 255.255.255.000, the network portion of the flow's IP address would be either 192.085.001 or 010.000.001 through 010.000.003.

Cyclic flows on TeraMetrics or TeraMetrics-based ports can use a bitmask instead of the subnet mask for the IP address if the value in the *Bitmask size* field (on the *Traffic* tab) is greater than the subnet mask size. Using the same example, if the bit mask size is 24 bits, the masked portion of the flow's address would be 10.0.1.x. *Figure 7-17* shows the relationship between an IPv4 address and the setup on the *IPv4 Networks* tab.

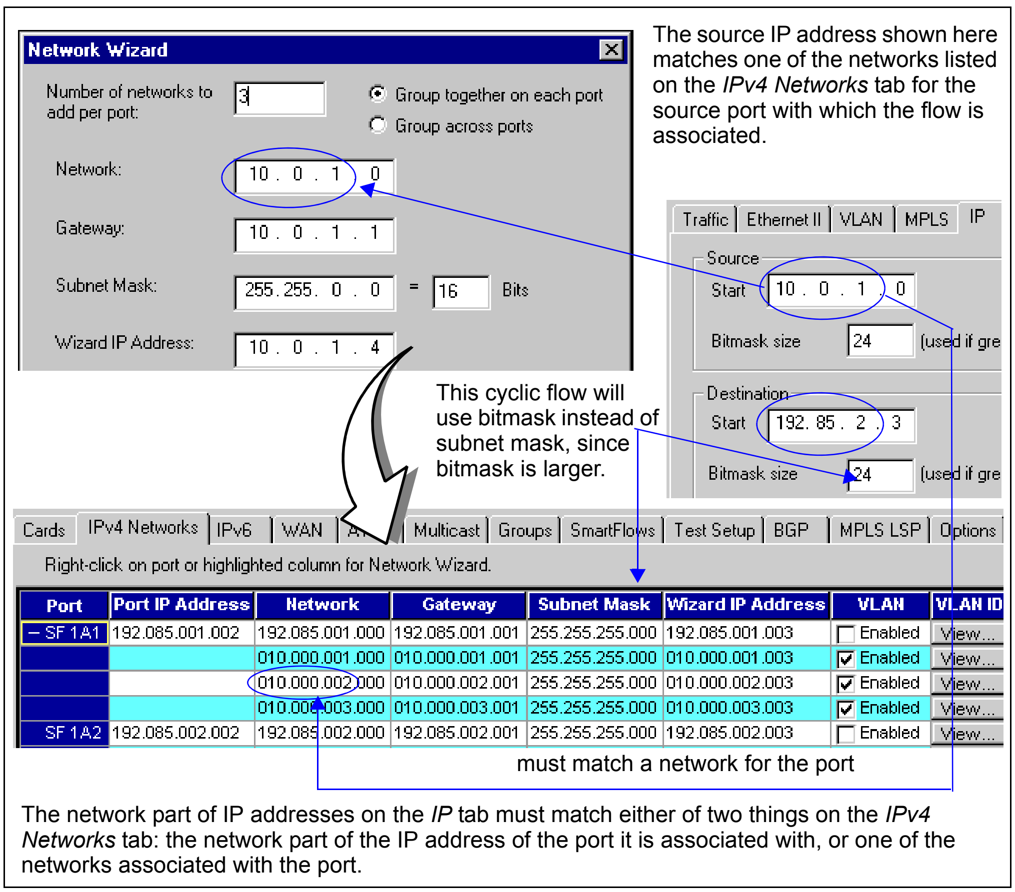

Figure 7-17. How IPv4 Flow IP Addresses Relate to the IPv4 Network Setup

Both source and destination IP addresses for a flow must be on the same network as the source and destination ports with which they are associated. If addresses do not match, when you click another tab, a message displays, prompting you to select an IP address associated with the port from the drop-down list. (See the example in *Figure 7-18*.) The list displays only addresses that are associated with the port used in the flow, so any one you select will be valid.

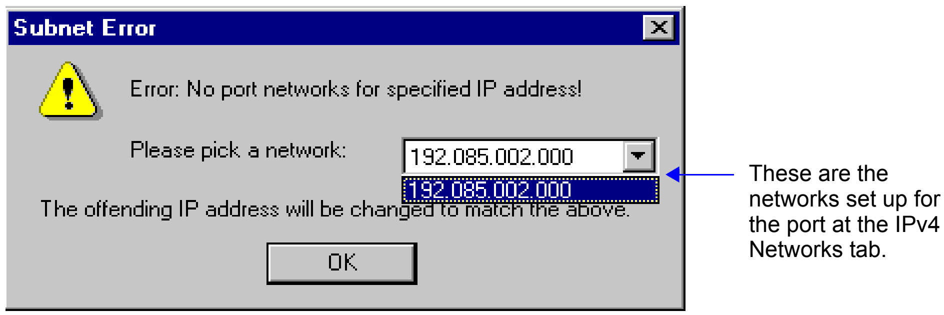

Figure 7-18. Error Message Box for Unmatching IPv4 Addresses

How SmartFlow Varies IP Addresses

To create more realistic test traffic, SmartFlow can vary IP addresses in test frames by incrementing them each time a frame is sent. SmartFlow increments IP addresses when you select the *SRC IP address* or *DST IP address* checkbox on the *Traffic* tab as the variable fields. Doing this enables the *Variable step* field on the *IP* tab.

You use the following fields to set up address variations:

- The *Variable count* field on the *Traffic* tab defines the number of variations to generate for the address field.
- The *Variable step* field on the *IP* tab specifies how much to increment the address field each time it is varied.
- With IPv4 flows on TeraMetrics and TeraMetrics-based modules, the *Bitmask size* field value on the *IP* tab enables you to specify how many bits in the address SmartFlow should vary, starting with the right-most bit.
- With IPv6 flows, the *Offset (bits)* field value specifies where in the address SmartFlow should vary the address bits within a default four-byte pattern. The pattern starts with the offset bit and moving leftward toward the most-significant bit in the address.

Note: For information about how SmartFlow calculates default IP addresses, see *"How SmartFlow Computes Default IP Addresses for Flows and Ports" on page 51*.

The best way to understand IP address variation is through examples.

- For IPv4 examples: *"IPv4 Address Variation Examples" on page 196*.
- For IPv6 examples: *"IPv6 Address Variation Examples" on page 200*.

IPv4 Address Variation Examples

With IPv4 traffic, addresses increment based on the network mask. SmartFlow does not increment the network portion of the address.

Note: For IPv4, do not specify address incrementing that will result in the IP broadcast address (all ones binary) in the host portion (netmask 255.255.255.255) or the network address (all zeros binary) in the network portion of the address. An error message displays if either condition occurs. For the steps to skip subnet or broadcast IP addresses, see *"Skipping Subnet and Broadcast IPv4 Addresses" on page 199*.

Table 7-3 shows examples of incrementing IPv4 addresses.

Table 7-3. Examples of Varying IP Addresses in IPv4

If the Starting IP address is...	and the Subnet Mask is...	and the Variable Count value is...	and the Variable Step is...	the sequence of IP addresses will be...
192.85.1.3	255.255.255.0	2	1	192.85.1.3 192.85.1.4 192.85.1.3
192.85.1.3	255.255.255.0	2	2	192.85.1.3 192.85.1.5 192.85.1.3
192.85.1.3	255.255.255.0	4	2	192.85.1.3 192.85.1.5 192.85.1.7 192.85.1.9 192.85.1.3
192.85.1.3	255.255.255.0	5	2	192.85.1.3 192.85.1.5 192.85.1.7 192.85.1.9 192.85.1.11 192.85.1.3
192.85.255.240	255.255.0.0	5	3	192.85.255.240 192.85.255.243 192.85.255.246 192.85.255.249 192.85.255.252 192.85.255.240 ERROR MESSAGE

Using Bit Masks When Incrementing IPv4 Addresses

With IPv4 cyclic flows transmitted from TeraMetrics and TeraMetrics-based ports, you can use bit masks to define what portion of the address does *not* increment. This then defines what bits in the address do increment. As with the subnet mask, the bit mask prevents a portion of an IPv4 address from incrementing. You can mask an area larger than the network portion of the address, including part of the host bits in the source and destination address of the flow.

Figure 7-19 shows the relationship of the subnet mask to the bit mask.

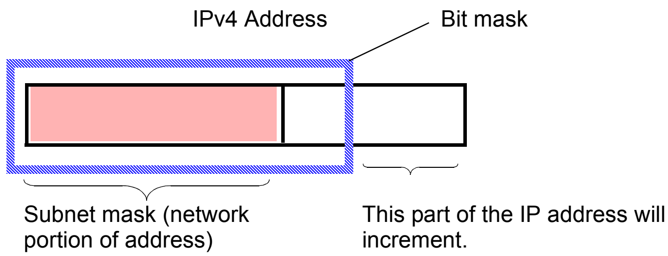

Figure 7-19. How the Subnet Mask Relates to the Bit Mask

Tip: By specifying a *different* bit mask size for the source address and the destination address, you can in effect have two different *Variable count* values (provided that SmartFlow had to adjust one of the values to maintain the bit mask): one for source and one for destination. This gives you more flexibility and control over address destination.

The bit mask starts at the most significant bit of the address and masks towards the least significant bit (left to right). (Refer to *Figure 7-19*.)

Some examples follow in order to illustrate the use of this option.

Example 1
Bit mask is smaller than subnet mask. Subnet mask is used.

Let us say you specify these parameters, where the bit mask size is *smaller* than the subnet mask size:

Subnet mask: 24 bits - 255.255.255.0
Variable count: 6
Bit mask size: 16
Variable step: Always 1 for IPv4 cyclic flows

SmartFlow will generate the following sequence of IP addresses (shown in both decimal and binary). The **bold** part of the binary addresses is the masked area, which in this case is from the subnet mask.

```
192.168.85.13   11000000.10101010.01010000.00001101 (base)
192.168.85.14   11000000.10101010.01010000.00001110
192.168.85.15   11000000.10101010.01010000.00001111
192.168.85.16   11000000.10101010.01010000.00010000
192.168.85.17   11000000.10101010.01010000.00010001
192.168.85.18   11000000.10101010.01010000.00010010
192.168.85.13   11000000.10101010.01010000.00001101
```

Example 2
Bit mask is larger than subnet mask. Bit mask is used. *Variable count* is adjusted.

Let us say you specify these parameters, where the bit mask size is *larger* than the subnet mask size:

Subnet mask: 24 bits
Variable count: 6
Bit mask size: 28
Variable step: Always 1 for IPv4 cyclic flows

SmartFlow will adjust the value for the *Variable count* from the specified value of **6** to **3** in order to maintain the bit mask and not increment that portion of the address.

The following sequence of IP addresses, shown in both decimal and binary, will be generated. The **bold** portion of the binary address shows the masked area, which in this case is from the bit mask.

```
192.168.80.13   11000000.10101000.01010000.00001101 (base)
192.168.80.14   11000000.10101000.01010000.00001110
192.168.80.15   11000000.10101000.01010000.00001111
192.168.80.13   11000000.10101000.01010000.00001101
192.168.80.14   11000000.10101000.01010000.00001110
192.168.80.15   11000000.10101000.01010000.00001111
```

Skipping Subnet and Broadcast IPv4 Addresses

When using incrementing addresses in IPv4 cyclic flows transmitted from a TeraMetrics or TeraMetrics-based port, you can skip the subnet and broadcast addresses. Select the *Skip subnet and broadcast* checkbox on the *Traffic* tab.

Example 1
Skip subnet and broadcast checkbox is *not* selected. Let us say you these specify parameters:

Subnet mask: 24 bits - 255.255.255.0
Variable count: 3
Variable step: Always 1 for IPv4 cyclic flows

With the *Skip subnet and broadcast* checkbox clear, the following sequence of IP addresses, shown in both decimal and binary, will be generated.

`192.168.80.1311000000.10101000.01010000.00001101 (base)`

`192.168.80.1411000000.10101000.01010000.00001110`

`192.168.80.1511000000.10101000.01010000.0000`**`1111`**

`192.168.80.1311000000.10101000.01010000.00001101`

`192.168.80.1411000000.10101000.01010000.00001110`

`192.168.80.1511000000.10101000.01010000.0000`**`1111`**

Notice that the broadcast address is included even though it contains all **1**s in the host portion of the address.

`192.168.80.15  11000000.10101000.01010000.0000`**`1111`**

Example 2
Skip subnet and broadcast checkbox is selected. Let us say you these specify parameters:

Subnet mask: 28 bits
Variable count: 3
Variable step: Always 1 for IPv4 cyclic flows

With *Skip subnet and broadcast* checkbox selected, the following sequence of IP addresses, shown in both decimal and binary, will be generated.

`192.168.80.13  11000000.10101000.01010000.00001101 (base)`

`192.168.80.14  11000000.10101000.01010000.00001110`

`192.168.80.13  11000000.10101000.01010000.00001101`

`192.168.80.14  11000000.10101000.01010000.00001110`

The address below is skipped because it represents the broadcast address (with all **1**s in the host portion of the address).

`192.168.80.15  11000000.10101000.01010000.0000`**`1111`**

IPv6 Address Variation Examples

With IPv6 flows, you use the *Offset (bits)* field on the *IP* tab to specify what address bits should increment. The offset sets the starting point of a four-byte pattern. SmartFlow increments the bits in the pattern by the amount set in *Variable step* field as test frames are generated.

The four-byte pattern starts at the specified offset bit, counting from bit 0 (the least-significant bit) and masks toward the most-significant bit. By default, the offset is 0. SmartFlow increments starting with the least-significant bit in the address.

You can place the four-byte pattern in any part of the IPv6 address, including both network address bits and host address bits.

The maximum step of 255 is relative to the byte boundary where the offset begins. Incrementing must take place within one octet in the four-byte pattern. For example, if you set *Offset (bits)* to **7** (8th bit of the first octet), the only allowable value for *Variable Step* would be **1**. A larger value would cause incrementing to overflow into the second octet. In contrast, if the *Offset* is **0** (1st bit of the first octet), the *Variable Step* value can be as much as **255**. Incrementing would affect all bits of the first octet but would not overflow into the second octet.

Figure 7-20 is an example of the offset and four-byte pattern in an IPv6 address. In this example, the *Offset (bits)* field value is **24**. The four-byte pattern extends from bit 24 to bit 55 (inclusive).

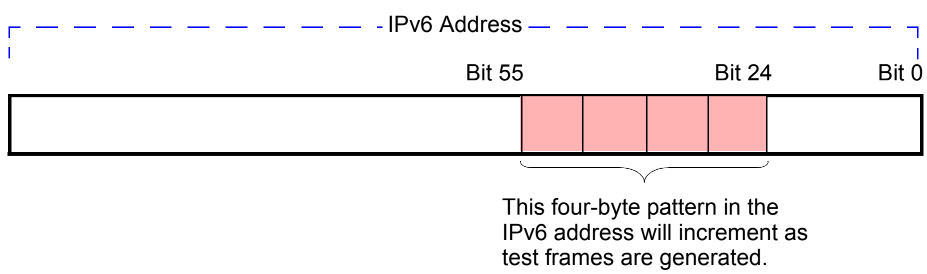

Figure 7-20. Using the Offset When Incrementing IPv6 Addresses

Important: With IPv6 flows, do not use an incrementing scheme that results in the first eight bits being FF:: in the address. This is a multicast address and SmartFlow currently does not support multicast in IPv6.

Table 7-4 on page 201 shows examples of incrementing IPv6 addresses.

Table 7-4. Examples of Varying IP Addresses in IPv6

If the Starting IP address is...	and the Variable Count value is...	and the Variable Step is...	...the sequence of IP addresses will be:
2001:0001:0000:0000:0000:C000:0103	2	1	2001:0001:0000:0000:0000:C000:0103 2001:0001:0000:0000:0000:C000:0104 2001:0001:0000:0000:0000:C000:0103
2001:0001:0000:0000:0000:C000:0108	8	1	2001:0001:0000:0000:0000:C000:0108 2001:0001:0000:0000:0000:C000:0109 2001:0001:0000:0000:0000:C000:010A 2001:0001:0000:0000:0000:C000:010B 2001:0001:0000:0000:0000:C000:010C 2001:0001:0000:0000:0000:C000:010D 2001:0001:0000:0000:0000:C000:010E 2001:0001:0000:0000:0000:C000:010F 2001:0001:0000:0000:0000:C000:0108
2001:0001:0000:0000:0000:C000:010F	2	1	2001:0001:0000:0000:0000:C000:010F 2001:0001:0000:0000:0000:C000:0110 2001:0001:0000:0000:0000:C000:010F
2001:0001:0000:0000:0000:C000:FFFE	3	2	2001:0001:0000:0000:0000:C000:FFFE 2001:0001:0000:0000:0000:C001:0000 2001:0001:0000:0000:0000:C001:0002 2001:0001:0000:0000:0000:C000:FFFE

Varying TCP/UDP Ports for a Flow

If you selected TCP or UDP for the *IP's next protocol* on the *Traffic* tab, click on the *TCP* or *UDP* tab to specify the source and destination ports for the flow. The TCP and UDP tabs contain the same fields. *Figure 7-21* shows the *UDP* tab.

If you chose *SRC Port* and/or *DST Port* as a variable field(s) on the *Traffic* tab, you can vary the port numbers. SmartFlow varies TCP and UDP ports by incrementing the decimal equivalent of the port set in the *Start* field. Incrementing the decimal number changes the type of port and thereby the type of traffic.

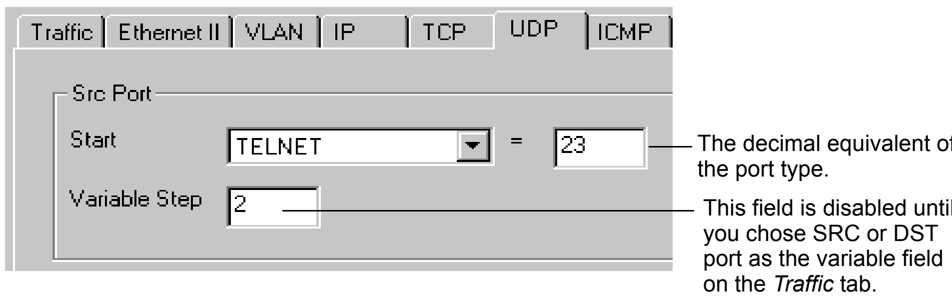

Figure 7-21. UDP Tab

For detailed information on fields on the *TCP* and *UDP* tabs, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Specifying ICMP Type for a Flow

If you selected ICMP (Internet Control Message Protocol) for the *IP's Next Protocol* on the *Traffic* tab, click on the *ICMP* tab to specify the type of ICMP message and code to associate with that message type for the flow. ICMP detects and generates messages about congestion, network errors, and timeouts. Many ICMP messages are identified by type, such as Echo (used by PING).

Figure 7-22 shows the *ICMP* tab.

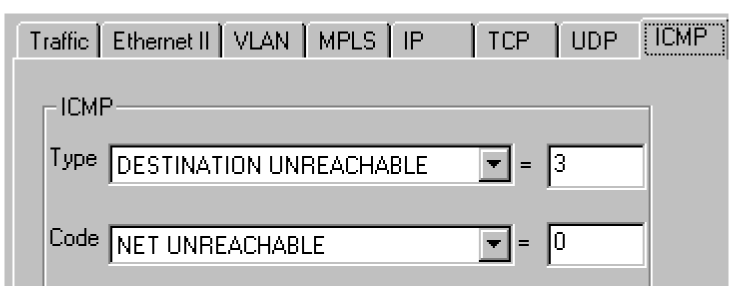

Figure 7-22. ICMP Tab

For detailed information on fields on the *ICMP* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

How to Set up (or Add) a Group Manually

Groups are used only as an aid to tracking test results. (Do not confuse *groups* as used here, which consist of SmartFlows, with IP Multicast groups.) During a test, flows are transmitted, not groups. You can use the Group Wizard to set up (or add) multiple flows in a group simultaneously or set up groups manually, one at a time. See *"Modifying Flows and Groups Created with the Wizard" on page 169* for information about how to set up flows and groups using the Group Wizard. *Figure 7-23* shows the *Groups* tab.

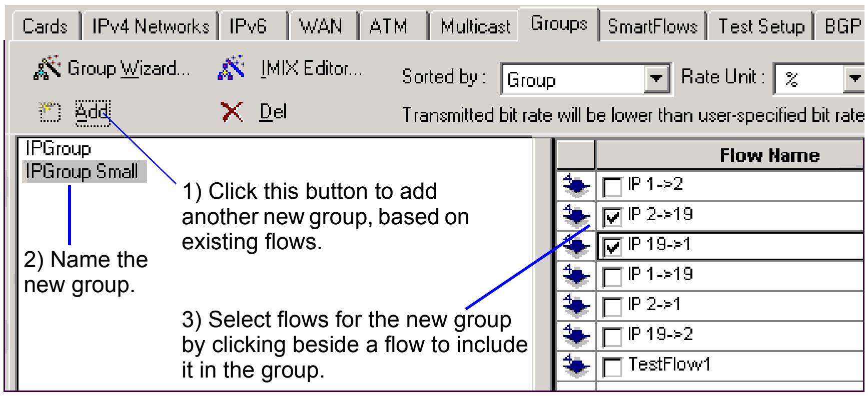

Figure 7-23. Groups tab

For detailed information on fields on the *Groups* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

The *Groups* tab lists all existing groups and flows that were created individually or with the Group Wizard. You can use the *Groups* tab to:

- Include or exclude selected groups in a test.
- View flows according to other sorting options, such as by VLAN or IP priority.
- Manually set up groups one at a time and assign existing flows to them.
- Manually set up a single group that contains all flows from a number of other groups.
- Modify the combination of flows assigned to groups.
- Access the *Group Wizard* to set up multiple flows in a group simultaneously.
- Assign a rate on a per flow basis. You can also assign a rate per flow basis based on its VLAN or CoS.
- Access the *IMIX Editor* to set up a group of flows with weighted packet sizes.
- Delete a group (and optionally all its flows).

Note: All the examples of sorting assume that flow names are the original names when the flows were created; they have not been modified.

The flow display varies according to the sorting option selected. Refer to the online Help for descriptions of each sorting option, field descriptions, and other information about the Group Wizard.

To change the VLAN priority or VLAN ID for a flow, click the **SmartFlows>VLAN** tab.

To change IP priority, click the **SmartFlows>IP** tab.

To use the Groups tab or manually set up a new group:

1. Select the sorting option by which you want to view flows from the **Sorted by** drop-down list.
2. Click the **Add** button to add a new group. It will be based on existing flows.
3. Enter a name for the new group and press **Enter**.
4. In the list of flows already created (right side of tab), click the box beside the flows that you want to include in the new group.
5. Repeat the above steps for each additional group that you want to create.
6. You can assign a specific rate for selected flows, as opposed to each flow sharing an even distribution of the port speed. To do this, select the rate unit from the **Rate Unit** drop-down list. Then enter a rate in the*%/Rate unit/Custom* field for the flow.

Note: When using a rate unit other than percentage (%), the bit rate that you specify is based on 100% of port load. This means that for any test load iteration less than 100% load, the actual transmission rate will be lower than the bit rate that you specify.

7. If you want to assign a specific rate to all flows associated with a particular VLAN or class of service, click the **CoS Wizard** button to access the **CoS Wizard**.
8. If you entered a specific rate for any flows, click the **Validate Rate** button. Validation is done according to transmitting port, and all flows on a port are validated together to make sure they do not oversubscribe the port speed.

To create a single group that contains all flows from a number of other groups:

1. Click the **Add** button on the **Groups** tab.
2. Enter a name for the new group.
3. Press the **Enter** key.
4. Right-click on the new group name.
5. Choose the **Include other groups....** option.
6. Select other listed groups to be included in the **Target group**.
7. Click **OK**. Flow IDs from all selected groups are copied to the target group.

Changing Which Flows Make up a Group

You can change which flows constitute a group without necessarily having to create new flows or delete flows altogether.

To modify the flow composition of a group:

1. In the list on the left side of the tab, highlight the group that you want to change. The flows that are currently in that group display a check in the checkbox beside the flow in the *Flow Name* column.
2. Check the boxes beside the flows that you want to add to the group, or uncheck the boxes beside the flows that you want to remove from the group.

Copying, Modifying, and Deleting Flows and Groups

Once you create flows (either manually or with the Group Wizard), you can copy them and then modify one or all of their attributes. You can quickly create a large number of flows with the Group Wizard, then manually modify only certain parameters. For more information, see *"Modifying Flows and Groups Created with the Wizard" on page 169*.

Copying Flows

Once you create one or more flows, you can copy them at the *SmartFlows* tab. Copying flows provides a fast way to create additional flows with the same properties. Then modify those properties for the new flows.

You can modify flows either individually and/or in multiples. You can copy and modify them in stages. For example, you may want to:

1. Copy a flow(s) with a low priority or traffic class.
2. Modify the priority or traffic class on the *IP* tab to a medium priority or class.
3. Copy those modified flows and modify them again to a higher priority/class or modify another attribute.

You can also use the right-click menu to modify certain attributes of multiple flows. (See *"Modifying Multiple Flows at Once" on page 207*.)

To copy flows:

1. Click the **SmartFlows** tab.
2. Highlight the flow(s) that you want to copy in the flow list on the left.
3. Click the **Copy Flow** button. The *Copy Flow* dialog box opens.
4. Enter the number of copies (up 1999) that you want to make of each highlighted flow.

Note: The number of copies of a flow(s) is limited to the maximum number of flows allowed for the port. (See *"Maximum Number of Flows per Port" on page 45*.) However, the maximum allowable number of copies allowed at a time is 1999. If you want a large number of copies of *multiple* flows, as long as each highlighted flow uses a different source port, you can make up to 1999 copies of each flow separately.

Figure 7-24 illustrates these steps.

To rename the flows, double-click on the name and type the new name.

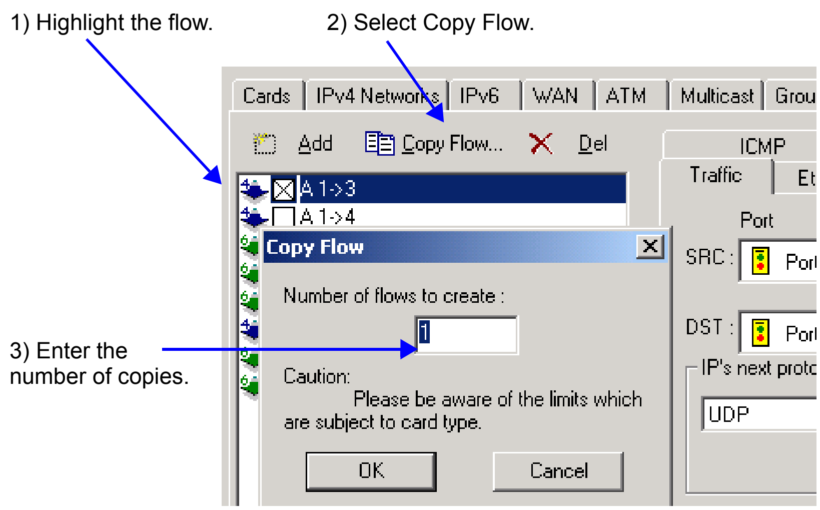

Figure 7-24. How to Copy a Flow

When you copy a flow, all the copies contain the same properties as the original flow. The default name of the first copy consists of the words *Copy of* prepended to the original flow name, in the format *Copy of x*, where *x* is the original flow name.

Example: The copy of flow **QoS_Low** is named **Copy of QoS_Low**.

You can also make multiple copies of a flow as well as copy multiple flows. Refer to the online Help for information about the default names for the copied flows.

Modifying Multiple Flows at Once

Once you create flows (either manually or by using the Group Wizard), you can copy them and then modify individual attributes for multiple flows at the same time. You manually modify certain attributes in multiple flows simultaneously at the *SmartFlows* tab.

You can modify flows individually and/or in multiples. You can copy and modify them in stages. You can modify the field in multiple flows by either:

- Copying the changed field value to each selected flow.
- Incrementing or decrementing the changed value in each selected flow.
- Generating (randomly) a value for the field.

You can modify the following attributes for multiple flows: Frame Length, Ethernet, IPv4, TCP, UDP, Variable Field, MPLS, IPv6 Flow, DHCP, Signature Fields, and Next Protocol. These attributes are listed in a right-click menu when you highlight multiple flows.

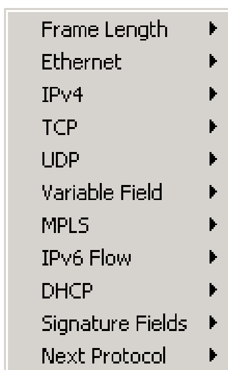

With each of these options, you can Copy Down, Fill Increment, Fill Decrement, or Fill Random.

The *Variable Field* and *Next Protocol* attributes allow only the *Copy Down* function.

Note: If the *Next Protocol* attribute is selected, the next protocol feature can be copied down to other flows. The next protocol parameter of the first selected flow is copied to all subsequent flows, where possible.

The *MPLS* and *IPv6 Flow* attributes enable or disable MPLS or IPv6, respectively, for all highlighted flows. Also, the extension header of the IPv6 flow can be copied down.

The *DHCP* attribute enables or disables the source and/or destination IP addresses.

The *Signature Fields* attribute can be turned on or off.

To modify multiple flows at once:

1. Click the **SmartFlows** tab.
2. Select a flow, and then select additional flows that you want to modify by using the **Shift** or **Ctrl** keys.
3. Right-click on any of the highlighted flows. SmartFlow displays a dynamic menu.

4 Move your cursor over the option that you want to modify. The menu changes dynamically, depending on your selection, as shown in *Figure 7-25 on page 208*.

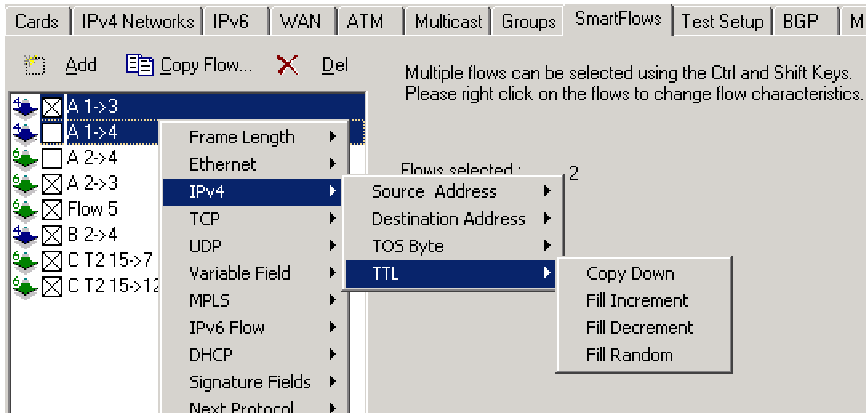

Figure 7-25. Dynamic Menu to Modify Multiple Flows

For more detailed information on modifying multiple flows and the functions that you can perform on them, refer to the online Help.

Deleting a Flow

Use the following steps to delete flows.

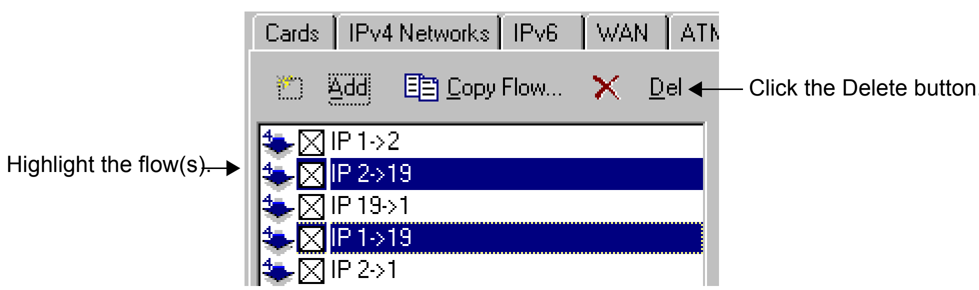

To delete a one or more flows:

1 Click the **SmartFlows** tab.

2 Highlight the flow(s) you want to delete in the list of flows.

3 Click the **Delete** button.
A confirmation box displays.

Deleting a Group

When you delete a group, you can also choose to delete all the flows in the group.

To delete one or more groups:

1. Click the **Groups** tab.
2. Highlight the group(s) you want to delete in the list of flows.
3. Click the **Delete** button.

   A confirmation box asks if you want to delete all flows in the group(s). Clear this box if you want to delete only the selected group(s), without also deleting the flows within them.

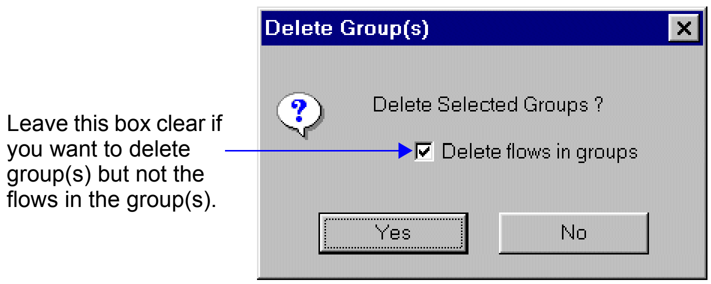

Important: You cannot delete individual flows from the *Groups* tab, only groups and the flows in them. If you highlight a flow in the flows list, then click the **Delete** button, SmartFlow deletes whatever group is highlighted in the list (and *all* its associated flows, if this option is selected).

Including and Excluding Flows for a Test

Before you run a test, verify at the *SmartFlows* tab that all of the flows that you want included in the test are enabled (selected) and those that you do not want are disabled (cleared).

Figure 7-26 illustrates flows that are enabled and disabled.

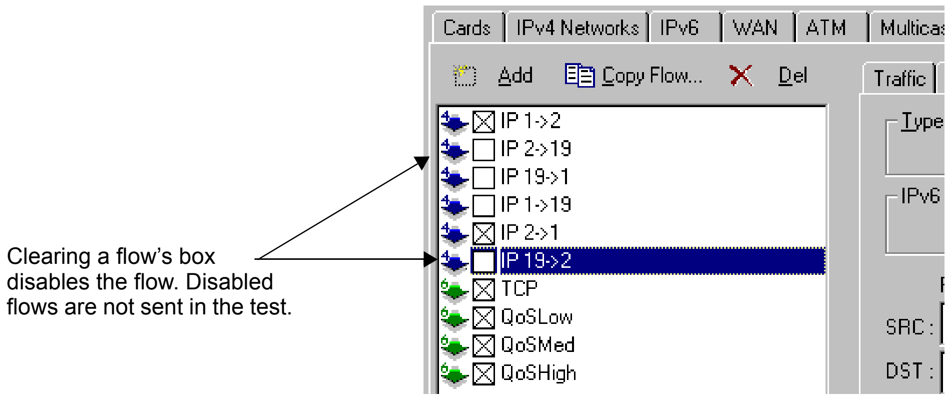

Figure 7-26. Enabling/Disabling Flows

Note: It is important to verify the flows before running the test. Even if a flow is included or excluded from a SmartFlow group, it may still be enabled or disabled on the *SmartFlows* tab.

Chapter 8

Set up and Run Tests

This chapter provides information about how to set up all tests as well as information specific to the Latency Distribution and Throughput tests, which require some additional parameters. It also explains how to start and stop tests.

In this chapter...

Test Setup Guidelines for Large Tests

Various factors contribute to a test's size and how many results are generated. SmartFlow stores all detailed test results in a Microsoft Access database, which has a 1 GB maximum size limit. If you want to run large tests and/or tests of long duration, test results may reach the 1 GB limit before the test finishes. If the 1 GB limit is reached prior to test completion, you receive this message:

```
SmartFlow database has reached maximum size.
```

If you receive this message, the test has aborted. In order to run the test, you need to adjust your test setup and/or port configuration to reduce the information written into the database.

Factors that Affect the Size of Test Results

Once you understand the relationship between the factors that affect the size of test results, you can adjust one or more of them and still run a large and/or long duration test to completion.

These factors affect the size of both the test and the results it generates:

- Results sampling
- Packet rate in results
- Number of SmartFlows
- Number and duration of test load iterations
- Number of ports
- Test duration
- Number of tests run.

Note: Although not directly related to the amount of test results generated, the size of your PC's memory and disk space can also be a factor when trying to run large tests. For large tests, it is recommended that your hard drive has at least 1 GB of free space. See *"PC Requirements" on page 64* for more information.

You may want to set up a test that runs for a *long* time to study the DUT's behavior. You can achieve this either with many test iterations or fewer iterations but a longer duration for each test iteration. Since SmartFlow takes one sample of test data for each test iteration per flow, fewer iterations will lessen the amount of data written to the database.

If you run a *shorter* test, you probably want to take results samples more often. For example, if you run a test for 10 seconds and you only specify one sample per load, you will only get one sample, so in this case you should increase the number of samples.

See *"Ways to Adjust Your Test Setup" on page 213* for more information.

Ways to Adjust Your Test Setup

In addition to the *Test Setup* tab to set up tests, you can also change various options on the *Options* tab that will affect both database size and also test duration and performance.

Figure 8-1 shows the *Options>Results* tab fields that are related to sampling test results and whether or not to write the samples (instead of all test data) to the database. For detailed explanations of these fields, see *"Results Display and Sampling Options" on page 292*.

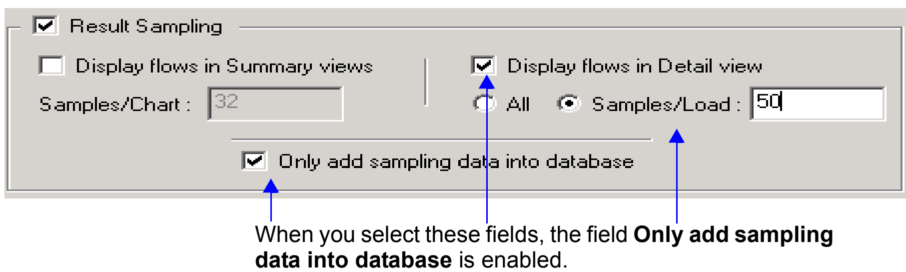

Figure 8-1. Results Sampling Fields on the Options Tab

Table 8-1 addresses how to adjust these factors to lessen the amount of data in test results. Lowering one factor may allow you to increase another that may be more vital to your test. For example, if it is important to have a large number of ports, perhaps you can increase the number of ports by using sampling to reduce the amount of data written to the database.

Table 8-1. Factors that Affect the Amount of Test Results and Database Size

Factor	Description / How to Adjust
Use of results sampling	*Overall effect:* Lessens the amount of results. Fewer samples per load iteration equals faster performance and smaller test results. Use of sampling restricts the amount of data being written to the database. SmartFlow only writes the data captured in the sample to the database. You can specify the number of test result samples that SmartFlow takes and writes to the database. On the *Options>Results* tab, do the following: **1** Select the **Display flows in Detail view** checkbox. **2** Instead of using the **All** option, enter a number of samples per load. The smaller the number, the fewer times SmartFlow samples results and writes them to the database. **3** Select the **Only add sampling data into database** checkbox. This option is enabled only if you select the *Display flows in Detail view* checkbox and click the *Samples/Load* button to exclude any flow data that is not included in the sampling. This can be done in addition to or instead of using the *Display flows in Summary views* checkbox and its related field.

Table 8-1. Factors that Affect the Amount of Test Results and Database Size (continued)

Factor	Description / How to Adjust
Including packet rate in results	*Overall effect:* Increases the amount of results. However, a larger sampling interval lessens amount of results. At regular sampling intervals, displays the number of packets transmitted and received for each port in the test. At the *Options>General* tab, you can either: • Clear the **Display Tx/Rx packet rate per port** checkbox in the *Packet rate* pane. • Check the **Display Tx/Rx packet rate per port** checkbox, but increase the value in the **Sampling interval (Sec)** field. A larger sampling interval time reduces the number of times that SmartFlow samples results and writes them to the database.
Number of ports	*Overall effect:* More ports increases the amount of Latency Over Time results while fewer ports lessens the amount of results. For other tests, only the number of flows affects the amount of test results generated. Either remove a port from the chassis or change the *Read State* column value (on the *Cards* tab) to *Inactive*.
Number of SmartFlows in test	*Overall effect:* Except for the Latency Over Time test, more SmartFlows increases the amount of results while fewer flows lessens the amount of results. Each individual SmartFlow utilizes one hardware stream. You can do any or all of these actions: • Make some flows inactive by clearing the checkbox beside the flow name on the *SmartFlows* tab. • Change the traffic pattern that you use to create flows. For example, the *Pair* traffic pattern produces a smaller number of flows than the *Fully-meshed* one. • Use cyclic SmartFlows instead of multiple individual SmartFlows. Cyclic flows use only one hardware stream and thus are only tracked once in test results. Check the **Cyclic** checkbox in the *Variable field(s) within flow* pane on the *SmartFlows>Traffic* tab or click the **Cyclic SmartFlows** button in the *SmartFlows* pane on the *Group Wizard-Multiple Flows* page.

Table 8-1. Factors that Affect the Amount of Test Results and Database Size (continued)

Factor	Description / How to Adjust
Number of test iterations	*Overall effect:* More iterations increases the amount of results while fewer iterations lessens the amount of results. The number of test iterations is determined by either: • Minimum and maximum test load specified together with the step size by which to increase the load on the *Test Setup>Test Iterations* tab. • For custom test loads, the value in the *# of iterations used in test* field on the *Custom Test Loads per Port (% Utilization)* dialog box accessed from the *Custom Loads Per Port Table* button on the *Test Setup>Test Iterations* tab.
Iteration duration	*Overall effect:* Longer duration increases the amount of results for the Latency Over Time test only. For the Latency SnapShot test, if the value in the *Capture frames per flow* field on the *Test Setup>Individual Tests* tab is large, it increases the amount of results. For other tests, unless you elected to display the packet rate in results, there is no effect on the database. You can do one or both of these actions: • On the **Test Setup>Test Iterations** tab, change the value in the **Duration Time (Sec.)** field to a lower number. • On the **Test Setup>Individual Tests** tab, lower the value in the **Capture frames per flow** field.
Total test duration	*Overall effect:* The longer a test takes to run, the more data is collected and may be written to the database. Total test duration is determined by the number of ports, number of flows, iteration duration, and number of iterations.
Number of tests run	*Overall effect:* More tests equal more test results. For example, the Jumbo test actually is a combination of the Latency, Latency Distribution, and Frame Loss tests.

Setting Up Test Parameters

You define general test parameters such as test duration, as well as ones specific to certain tests (Latency Distribution and Throughput), at the *Test Setup* tab. To access the *Test Setup* tab, click on the tab from the main window. The values that you set here will apply to each test run until modified. Therefore, it is best to verify these parameters before each test. For example, if you enter a value of 15 in the *Duration* field, then run the Frame Loss test and then the Latency test, each iteration of each test will run for 15 seconds.

For information about setting the minimum and maximum test rates for tests using Gigabit and POS cards, see *"Minimum/Maximum Rates for Gigabit and POS Ports" on page 287*.

The *Test Setup* tab includes these subtabs:

- *Test Iterations*
  Use this tab to specify general test parameters that apply to every test, such as test loads and duration. However, if you run a Throughput test, the rate parameters specified for throughput are used instead of the values in the *Iterating across traffic load* fields. See *"Test Iterations Tab" on page 217* for more information.
- *Learning*
  Use this tab to specify learning frame parameters such as size and rate. Learning frames are sent prior to the start of the test in order to register the MAC addresses of each flow involved in the test. If ARPs are enabled (on the *Cards* tab), SmartFlow does not send out learning packets in addition to ARPs. See *"Learning Tab" on page 219* for more information.
- *Individual Tests*
  Use this tab to specify parameters for the Latency, Latency Over Time, Throughput, and Latency SnapShot tests. See *"Individual Tests Tab" on page 220* for more information.
- *SmartTracker*
  Use this tab to specify parameters for the SmartTracker test. See *"SmartTracker Tab" on page 221* and *Chapter 16, "SmartTracker Test."*
- *Sample Iteration*
  Use this tab to specify parameters for running a sample iteration of your test, including the load and duration. You can also select to confirm receipt of an iteration packet per receiving port (and per flow for multicast traffic). SmartFlow discards results after the sample iteration finishes. See *"Sample Iteration Tab" on page 222* for more information.
- *DHCP*
  Use this tab to customize the DHCP port timeout, port retries, request rate values, and lease expiry time. See *"DHCP Tab" on page 224* for more information.

Test Iterations Tab

The *Test Iterations* tab contains fields that apply to *every port* in every test. However, if you run a Throughput test, the rate parameters specified for throughput are used instead of the values in the *Iterating across traffic load* fields. *Figure 8-2* shows the *Test Iterations* tab.

The values that you specify here apply to each test run until you modify the values. It is best to verify these parameters before each test.

To set rates on a *per port* basis (instead of or in addition to custom test loads), use the *Groups* tab once you have set up flows. For more information, see *"Rates per Flow" on page 264*.

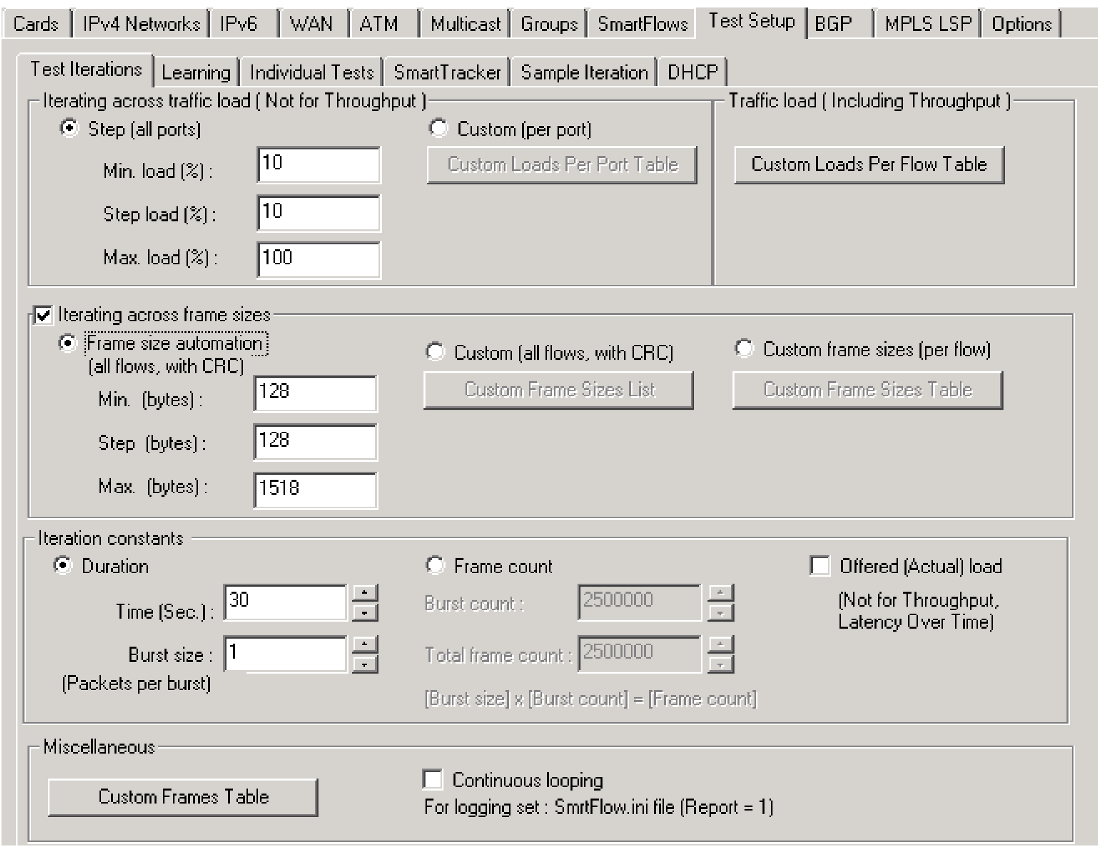

Figure 8-2. Test Iterations Tab

For information about how to use these fields to control the size of test results, see *"Ways to Adjust Your Test Setup" on page 213*.

For information about frame size automation, see *"Using Frame Size Automation (Global for All Flows)" on page 236*. For information about custom frame sizes, see *"Using Custom Frame Sizes per Flow" on page 228*.

If you specify a rate in the*%/Rate unit/Custom* field on the *Groups* tab for the flow, it will be a percentage of the rates specified in the *Iterating across traffic load* pane of the *Test Iterations* tab. For more information, see *"Rates per Flow" on page 264*.

For a description of each field on the *Test Iterations* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

To allow results from continuous looping to display in the .csv file:

1 Open the `SmrtFlow.ini` file in a text editor.

2 Under the *General* section, change value for the parameter called **Report** from **0** to **1**.

Figure 8-3 is an example of the results that appear in the `report.csv` file for a continuously looping test.

==========								
Step = 2	Time: 14:23:52 10/17/2001							
Name	Flow Load	Stream Cnt	Packet Cr	Received	Duplicated	Lost	Sequenc	Framesize
All	N/A	4	1520256	0	0	1520256	0	N/A
A DIFF0	N/A	1	380064	0	0	380064	0	N/A
A DIFF2	N/A	1	380064	0	0	380064	0	N/A
A DIFF4	N/A	1	380064	0	0	380064	0	N/A
A DIFF6	N/A	1	380064	0	0	380064	0	N/A
A D0 3:2A	15	1	380064	0	0	380064	0	128
A D2 3:2A	15	1	380064	0	0	380064	0	128
A D4 3:2A	15	1	380064	0	0	380064	0	128
A D6 3:2A	15	1	380064	0	0	380064	0	128
==========								
Step = 1	Time: 14:24:05 10/17/2001							
Name	Flow Load	Stream Cnt	Packet Cr	Received	Duplicated	Lost	Sequenc	Framesize
All	N/A	4	253368	0	0	253368	0	N/A
A DIFF0	N/A	1	63342	0	0	63342	0	N/A
A DIFF2	N/A	1	63342	0	0	63342	0	N/A
A DIFF4	N/A	1	63342	0	0	63342	0	N/A
A DIFF6	N/A	1	63342	0	0	63342	0	N/A
A D0 3:2A	2.5	1	63342	0	0	63342	0	128
A D2 3:2A	2.5	1	63342	0	0	63342	0	128
A D4 3:2A	2.5	1	63342	0	0	63342	0	128
A D6 3:2A	2.5	1	63342	0	0	63342	0	128
==========								
Step = 2	Time: 14:24:18 10/17/2001							

The test started over here. (→ Step = 1)

Figure 8-3. Results for a Continuously Looping Test

The report.csv file displays information equivalent to the *Detail* tab of test results. The *Stream Count* column refers to the number of streams in the flow.

Note: If you run another test, the results in this file will be overwritten.

Learning Tab

You can use the *Learning* tab to specify how learning frames should be sent to the DUT to register the MAC addresses of each flow involved in the test. For more information about learning, see *"Address Resolution" on page 53*.

If ARPs are enabled, SmartFlow sends out ARPs during the learning phase.

For cyclic flows transmitted from LAN-3201x ports: SmartFlow will send an ARP request for each source IP address (through incrementing) in the cyclic flow.

Note: Learning adds to the number of streams that need to be built for a test. Each card has a maximum number of streams, so if you plan on having a large number of flows per port (or card), it is recommended that you estimate the number of actual streams for the ports involved in the test to make sure their number does not exceed the card's maximum. See *"Maximum Number of Flows per Port" on page 45* for more information.

Important: When setting up a 32K stream test on the receive port, you need to turn off learning in SmartFlow because the receive port does not support this many learning frames. Also, provide the DUT with an appropriate path for the flows before running the 32K receive stream test. Do this by configuring the DUT with the path or run a SmartFlow test with a small number of streams that will educate the DUT with the path on the receive side.

The *Learning* tab is shown in *Figure 8-4 on page 220*. For field descriptions, refer to the SmartFlow online Help. Press **F1** over the tab or select **Help** from the menu bar.

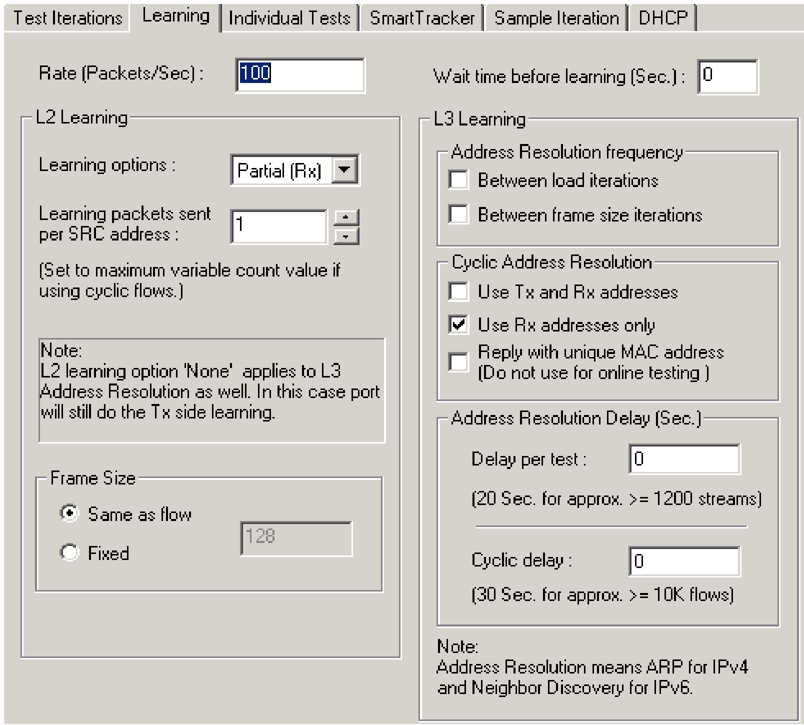

Figure 8-4. Learning Tab

Individual Tests Tab

The *Individual Tests* tab contains fields that apply only to the Latency Distribution, Latency Over Time, Latency SnapShot, and Throughput tests.

The *Latency options for Non-XD TeraMetrics-based modules* pane fields on the *Individual Tests* tab apply only to how latency for all TeraMetrics-based receiving ports appear in tests that measure latency. These tests include Jumbo, Latency, and Latency Over Time. For more information about these tests, refer to the appropriate test chapter.

For a description of each field on the *Individual Tests* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Figure 8-5 shows the *Individual Tests* tab.

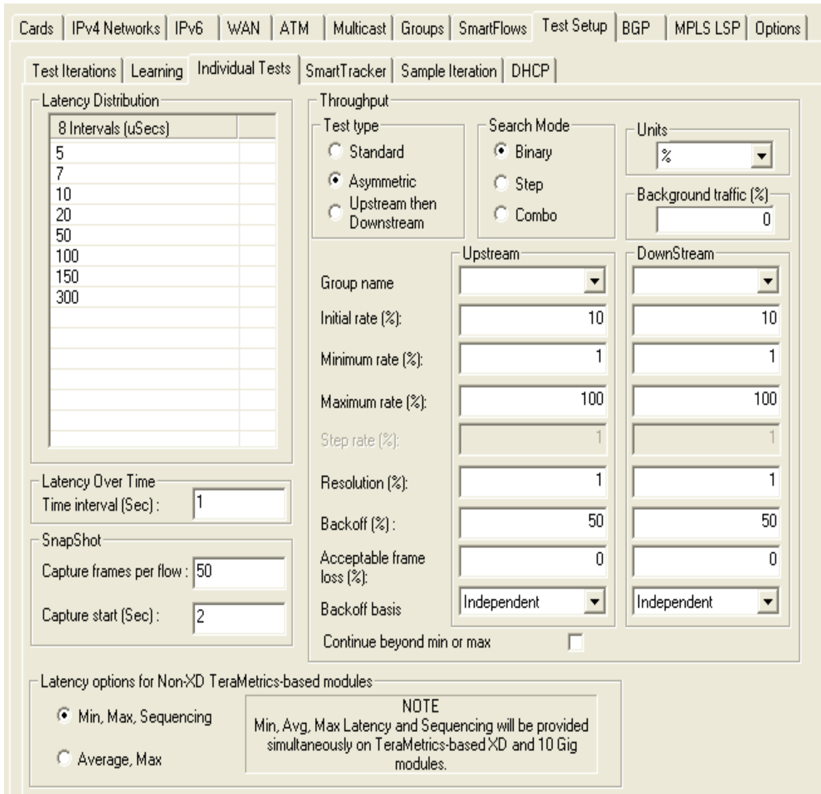

Figure 8-5. Individual Tests Tab

Note: The Throughput test rate fields override the rate fields defined in the *Iterating across traffic load* pane of the *Test Iterations* tab.

SmartTracker Tab

The *SmartTracker* tab contains fields that apply to the SmartTracker test.

For information about this test, see *Chapter 16, "SmartTracker Test"* or the SmartFlow online Help. For a description of each field on the *SmartTracker* tab, refer to the online Help. Press **F1** at the tab or select **Help** from the menu bar.

Figure 8-6 shows an example of a custom field being tracked with a custom traffic filter.

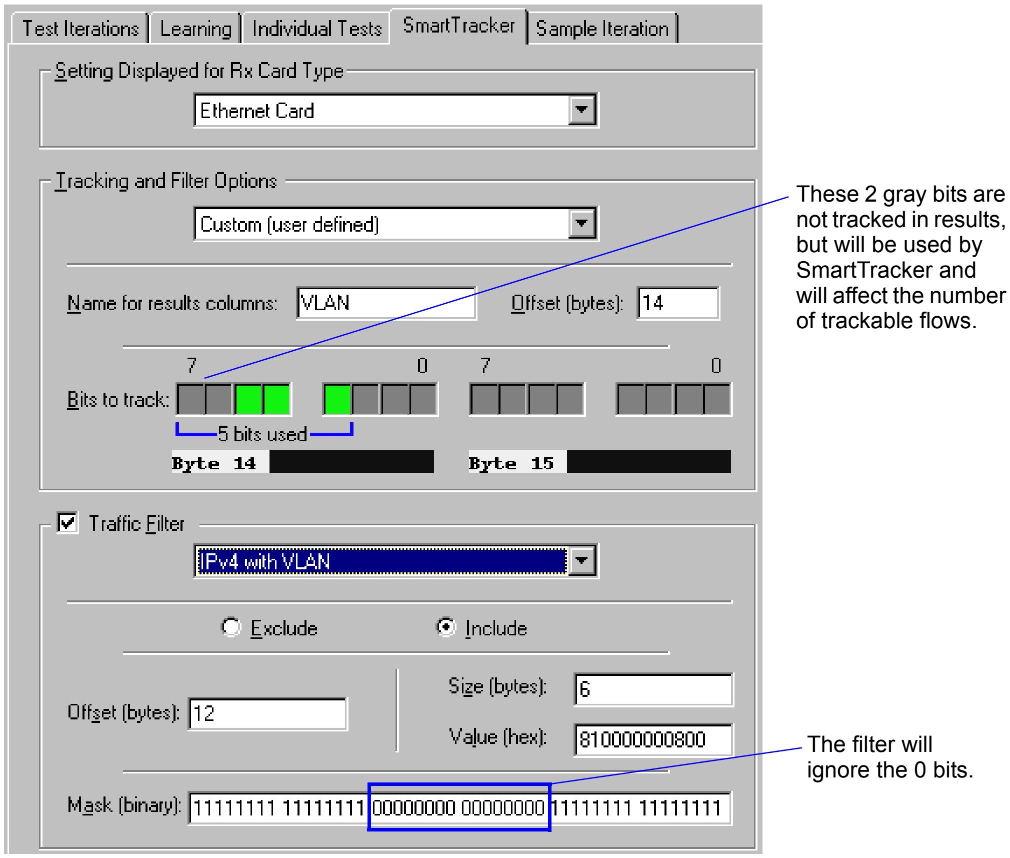

Figure 8-6. Test Setup SmartTracker Tab

Sample Iteration Tab

The *Sample Iteration* tab contains fields that define a sample test iteration. A sample iteration is a portion of the test that is run using your existing configuration of ports and flows. You specify the duration of the iteration and the load at the *Sample Iteration* tab, as shown in *Figure 8-7*.

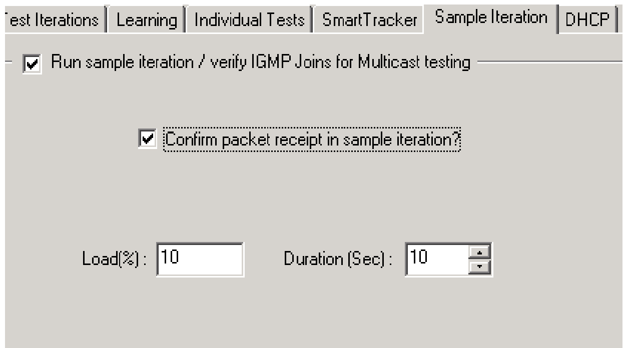

Figure 8-7. Sample Iteration tab

In a relatively short time, running a sample test iteration helps you:

- Determine if the test setup is correct and the test will run, without having to spend time running the actual test. This is especially useful if the actual test is of long duration.
- *For multicast and/or unicast traffic:*
  Pre-populate the DUT forwarding tables prior to the actual test traffic being sent. (This is also known as *fast-path learning*.)
- *For multicast traffic:*
  Verify that multicast hosts have successfully joined the group (if at least one packet is successfully received). The sample iteration is sent out after all multicast subscribers have sent out join requests and the IGMP join requests timer has elapsed.
- *For multicast traffic:*
  "Train" certain DUTs to form multicast associations (source IP addresses with multicast group addresses) before the test starts. This can serve as a type of learning phase or fast-path learning, since no learning packets or ARP requests are used with multicast traffic.

You can also confirm receipt of at least one packet at each receiving port (and also for each flow destined for the port, for multicast traffic) by selecting the *Confirm packet receipt in sample iteration?* checkbox. If the confirmation fails, a warning message displays prompting you to either:

- *Retry the sample iteration.*
  SmartFlow attempts to re-send sample test iteration packets.
- *Abort the sample iteration.*
  SmartFlow stops the test.
- *Ignore the non-receipt of a sample test iteration packet at any of the receiving ports.*
  SmartFlow starts sending actual test packets.

After the sample iteration finishes (and packet receipt is confirmed, if this option was selected), SmartFlow discards the sample results and starts transmitting actual test packets.

For a description of each field on the *Sample Iteration* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

DHCP Tab

The DHCP tab contains fields that determine the port timeout, port retries, request rate, and lease expiry time for DHCP as shown in *Figure 8-8*.

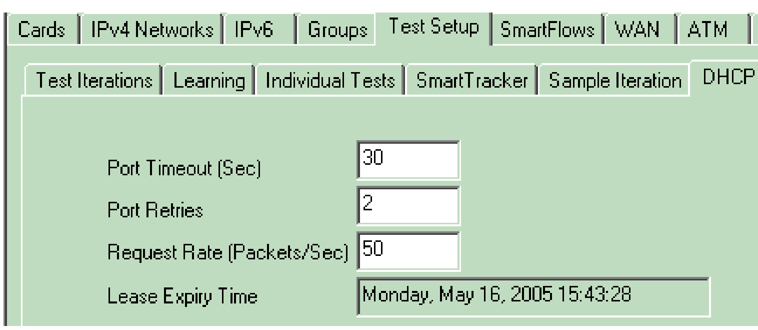

Figure 8-8. DHCP Tab

Dynamic Host Configuration Protocol (DHCP) enables individual computers on an IP network to extract their configurations from a server (the DHCP server) or servers; in particular, servers that have no exact information about the individual computers until the information is requested.

DHCP is integrated into SmartFlow as an additional step in the configuration before any test takes place.

For a description of each field on the *DHCP* tab, DHCP theory of operation, and procedures to use DHCP, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Varying Frame Sizes

You can use a constant frame size for all flows throughout the test, or you can vary the size of frames generated during a test. Varying frame sizes makes the test traffic more realistic, since in real world traffic a wide range of frame sizes can exist and in any order.

Note: You can also vary the test load to make traffic more realistic. See *"Varying Test Loads" on page 253*.

You can vary frame size in several ways:

- *Globally* in all flows after each complete set of load iterations is run.

  When you use the *Frame size automation (all flows, with CRC)* option on the *Test Iterations* tab, SmartFlow transmits one size of frames (for all flows) at each specified load, then increases or decreases frame size for all flows at each specified load. The change in size is accomplished either by:

  - Steps (an increase/decrease by this amount)
  - Custom frame size list.

  For more information, see *"Using Frame Size Automation (Global for All Flows)" on page 236*.

- *Per flow in each test iteration*.

  When you use the *Custom frame sizes (per flow)* option on the *Test Iterations* tab, SmartFlow transmits frames in *each flow* according to a predefined, custom sequence of frame sizes and/or random sizes. The True Random (R) option varies frame sizes within a single flow and iteration. This option best simulates real-world conditions in which a flow can contain user data of various frame sizes. Click the **Custom Frame Sizes Table** button to specify the actual frame sizes.

  For more information, see *"Using Custom Frame Sizes per Flow" on page 228*.

  Note: The *Custom frame sizes (per flow)* option does not apply to a Throughput test. The throughput cannot be determined because this option uses a set number of test iterations as well as frame sizes that change with each iteration.

- *Per flow with an IMIX set of frame sizes weighted by percentages.*

  When an IMIX set is used, SmartFlow creates a flow for each frame size in the IMIX set. Each IMIX flow transmits only that frame size, and the number of frames transmitted is within the ratio (percentage) specified in the IMIX set to the other frame sizes/flows in that set. In effect, this method is another way of setting the rate for the flow since each IMIX flow has a unique rate.

  For more information, see *"Using an IMIX Set of Frame Sizes for Flows" on page 240*.

All frame sizes displayed in SmartFlow always include CRC. Padding is never included.

Use *Table 8-2* to choose a method of frame size variation for your test.

Table 8-2. Summary of Methods to Vary Frame Size

If you want to ...	Then use ...
Vary frame size by test load iteration only (all flows at same size).	Frame size automation [with either the *Step (all ports)* or *Custom (per port)* option].
• Increase frame sizes by stepped amounts for all flows in each test load iteration.	Frame size automation with the *Step (all ports)* option.
• Vary frame sizes by custom list for all flows after each test load iteration.	Use frame size automation with the *Custom (per port)* option.
Vary frame size per flow by a custom list, after each test load iteration.	*Custom frame sizes (per flow)* option (with predefined sequence of sizes) and enter a sequence of *custom sizes.*
Vary frame size per flow by a randomly generated list of sizes, after each test load iteration.	*Custom frame sizes (per flow)* option with the *Fill Random* option of the right-click menu. Random sizes fall within default size ranges for the transmitting card/module technology and sizes display in the sequence (list).
Simulate an Internet mix or other frame size usage pattern with weighted frame sizes varied per flow.	*IMIX flow* option when creating flows with the Group Wizard. (First create an IMIX set with the *IMIX Editor*.) One flow is created for each frame size in the pattern. Each flow transmits the specified percentage of frames of that size in relation to the other flows in that IMIX set.
Vary frame size from frame to frame within a flow with truly random sizes generated at runtime, after each test load iteration.	*Custom frame sizes (per flow)* option and enter **"R"** in the sequence for that flow. **Note:** The following cards/modules support True Random (R) option: LAN-33xx, POS-3504/3505, POS-3510/3511, POS-3518/3519, XLW-3720/3721, and XFP-3730/3721.

Table 8-2. Summary of Methods to Vary Frame Size (continued)

If you want to ...	Then use ...
Vary frame size from frame to frame within a flow with pseudo-random sizes that are reproducible in future tests and are generated at runtime, after each test load iteration.	A non-zero value in the *Random seed* field on the *Options* tab. Use the *Custom frame sizes (per flow)* option and enter "**R**" in the sequence for that flow. Notes: 1) The following cards/modules support True Random (R) option: LAN-33xx, POS-3504/3505, POS-3510/3511, POS-3518/3519, XLW-3720/3721, and XFP-3730/3731. 2) Rates per flow on the *Groups* tab are not available when using the True Random (R) option for frame sizes.

Figure 8-9 shows the *Test Iterations* tab, where you specify the method to use to vary frame sizes.

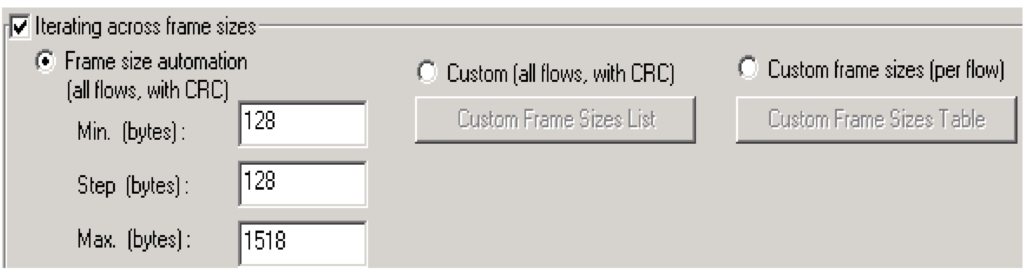

Figure 8-9. Frame Size Options on Test Iterations Tab

The *Frame size automation (all flows, with CRC)* option becomes disabled when you select one of the other two options in this pane.

The *Custom frame sizes (per flow)* option varies the frame size on a per flow basis.

The *Custom Frame Sizes Table* option specifies the actual frame sizes.

Using Custom Frame Sizes per Flow

The *Custom frame sizes (per flow)* option on the *Test Setup>Test Iterations* tab allows you to specify a different frame size per flow within the *same* test load iteration. In other words, each frame size corresponds to a test load iteration so the frame size *per flow* changes *each* time the traffic load changes. This makes test traffic more realistic than having uniform frame sizes for all flows.

Assign which sequence of sizes that you want for a flow on the *SmartFlows>Traffic* tab.

Note: You can specify a truly random ("R") frame size sequence for a flow. In this case, each *frame* in the flow will vary during the same iteration. For more information about random frame sizes, see *"Using Random Frame Sizes" on page 231* for more information.

Example:
This example shows a section of how a test using the *Custom frame sizes (per flow)* option on the *Test Setup>Test Iterations* tab runs. Some of the test setup parameters include:
Min. Load is 50%
Step Load is 10%
Seq. 1 starts with size 64, then 128
Seq. 2 starts with size 5058, then 8192

Port 1 transmits two flows, Flow A and Flow B. Flow A uses Seq.1 and Flow B uses Seq. 2. Port 1 transmits the frame sizes shown in *Table 8-3* for the specified test duration:

Table 8-3. Example of Custom Frame Sizes per Flow

Load	Flow	Frame Sizes
50%	Flow A Flow B	64, 64, 64, 64 5058, 5058, 5058, 5058
60%	Flow A Flow B	128, 128, 128, 128 8192, 8192, 8192

Notice how the same size frame per flow is transmitted during an iteration. The test continues changing frame sizes at each new test load iteration, according to the predefined sequence of sizes assigned to that flow. You can specify the sequences of frame sizes for flows to use during a test in the *Custom Frame Sizes Per Flow (bytes with CRC)* dialog box accessed by clicking the *Custom Frame Sizes Table* button. See *"Specifying a Custom Frame Size Sequence" on page 229* and *"Using Random Frame Sizes" on page 231* for more information.

Specifying a Custom Frame Size Sequence

The *Custom frame sizes (per flow)* option on the *Test Setup>Test Iterations* tab allows you to use a predefined series of frame sizes to use at each iteration of a test, known as a *sequence*. You can assign a sequence to one or more flows. You define these sequences in the *Custom Frame Sizes Per Flow (bytes with CRC)* dialog box accessed by clicking the *Custom Frame Sizes Table* button. (See *Figure 8-10*.)

Tip: If you want to have both VLAN tagged and non-tagged frame sizes with the *same size payload*, set up one sequence with sizes that include the VLAN tag (4 bytes more) and one sequence without VLAN tags (4 bytes less).

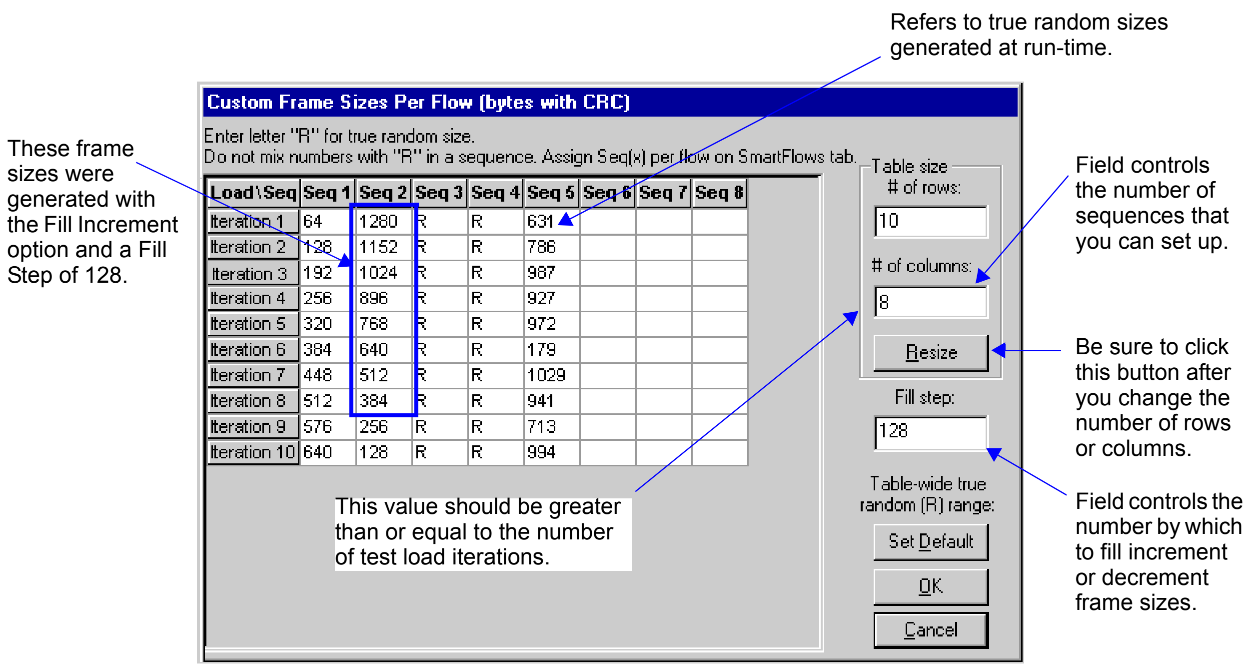

Figure 8-10. Setting up a Custom Frame Size Sequence

Each column must have the same number of iterations and must be either all filled or all blank. You can set up as many sequences as you want, depending on your computer's memory resources.

Important: If you change the number of rows or columns in the table, make sure you click the **Resize** button for the changes to take effect.

Keep in mind that frame sizes that you enter for a sequence include header information such as CRC and any VLAN tag or MPLS label, if applicable.

To access the Custom Frame Sizes Per Flow (bytes with CRC) dialog box:

1. From the **Test Setup** tab, click the **Test Iterations** tab.
2. Click the **Custom frame sizes (per flow)** button.
   If this option is disabled, select the **Iterating across frame sizes** checkbox.
3. Click the **Custom Frame Sizes Table** button.

For a description of each field on the *Custom Frame Sizes per Flow (bytes with CRC)* dialog box, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar. The online Help also explains how to determine the number of rows to set and how to add or change multiple cells.

When *Test Load* parameters are as shown in *Figure 8-11*, you know that there will be nine test load iterations. Therefore, you can set the *# of rows* field to nine or more.

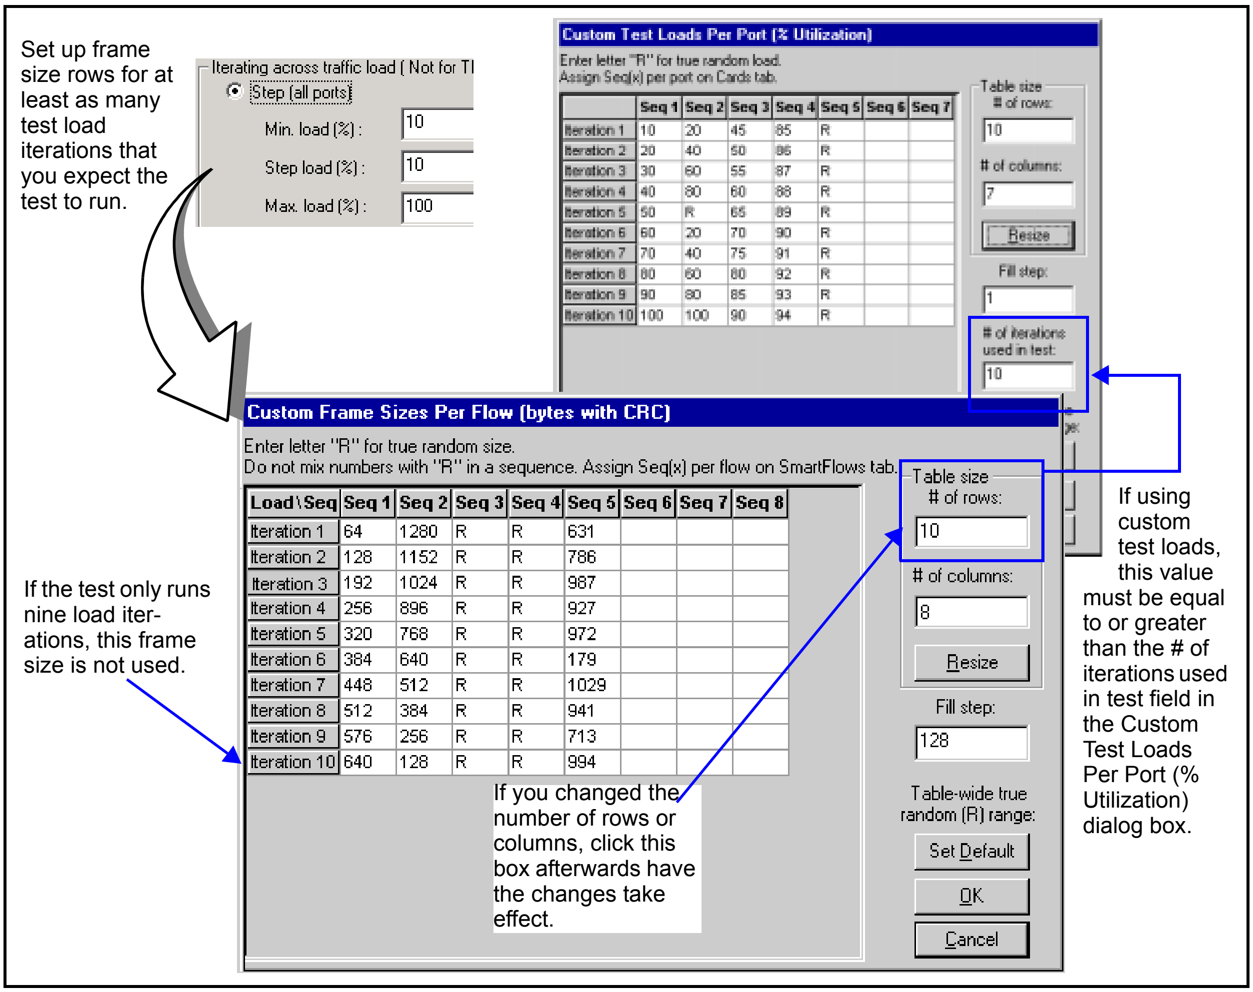

Figure 8-11. Setting the Number of Rows for the Custom Frame Size Table

Using Random Frame Sizes

SmartFlow provides two different ways that you can generate random frame sizes:

- *Fill Random* (from the right-click menu) – Fills table cells with random sizes.
  - Frame sizes are uniform with each flow for the specified duration of the test load iteration.
  - Frame sizes per flow change with each new test load iteration.
  - The actual sizes are displayed in the *Custom Frame Sizes Per Flow (bytes with CRC)* dialog box.
  - SmartFlow software generates random sizes with no pattern to them.
- *True Random (R)* – Entering R in a cell generates random sizes at runtime.
  - Frame sizes vary *from frame to frame* within each flow for the specified duration of the test load iteration.
  - Frame sizes change per frame within each flow with each new test load iteration.
  - The *Custom Frame Sizes Per Flow (bytes with CRC)* dialog box column displays an "R" instead of the random frame size because frame sizes are generated at runtime.
  - Card/module firmware generates random sizes with no pattern to them.
  - Frame sizes do not show up in results—zeroes show instead.

  The following cards/modules support the True Random (R) option: LAN-33xx, POS-3504/3505, POS-3510/3511, POS-3518/3519, XLW-3720/3721, and XFP-3730/3731.

When you use True Random frame sizes, SmartFlow calculates the utilization rate based on the maximum frame size specified for the random frame size range.

Note: When assigning flows to sequences at the *SmartFlows>Traffic* tab, sequences with Rs will not appear in the *Seq.* field's list of sequences unless the flow's transmitting card supports the true random (R) option.

Figure 8-12 on page 232 illustrates how true random (R) and fill random frame size generation methods differ.

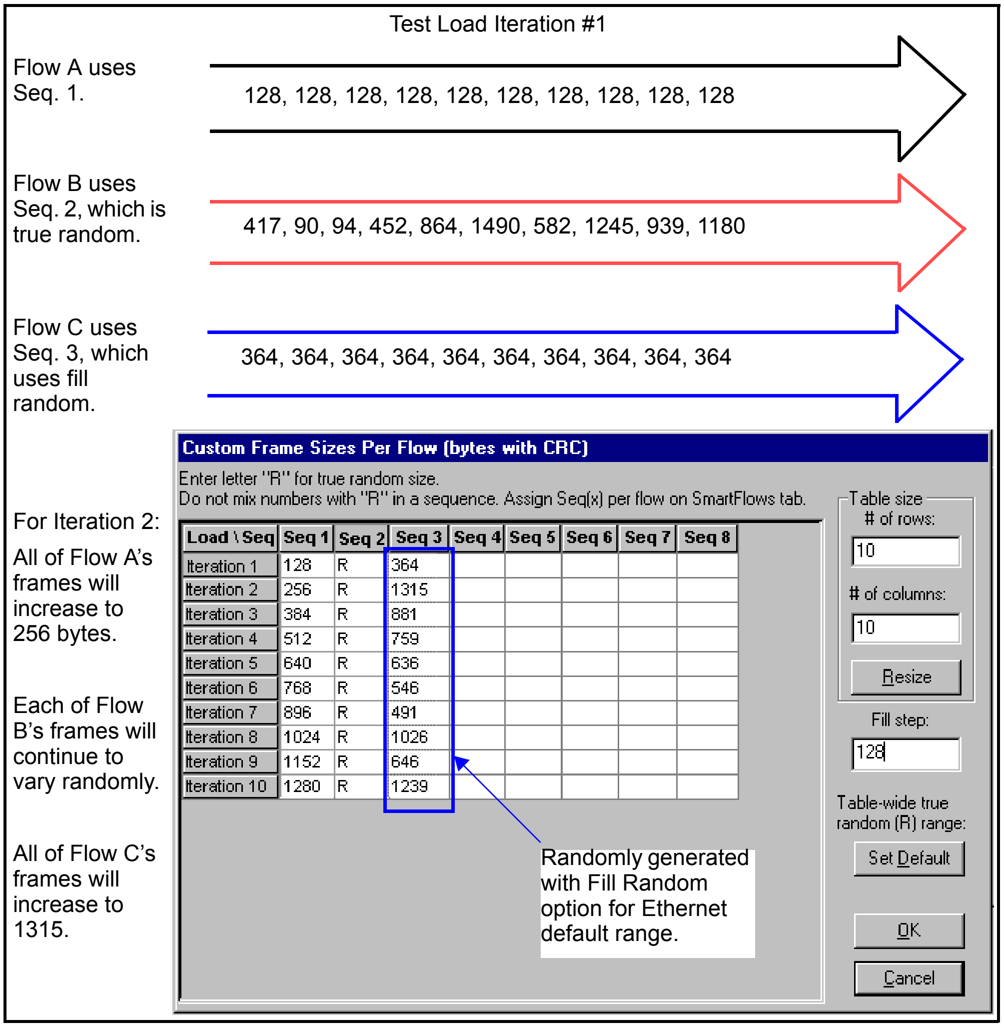

Figure 8-12. How Random Frame Size Generation Methods Differ

Notice how each frame in Flow B, which uses true random (R), varies in size even within the same iteration.

Tip: If you want to reproduce the test using the same set of random frame sizes, enter a value in the *Random seed* field on the *Options* tab. Before you rerun the test, make sure this same value appears in the field. For more information about this option, see *"Setting up Global Application Preferences" on page 288*.

For more information about card type and ranges, see *"Setting the Range for Random Frame Sizes" on page 233*.

Setting the Range for Random Frame Sizes

Both the fill random and true random (R) methods can generate frame sizes from within a specified size range. Default size ranges per technology (Ethernet, POS, WAN) are shown as guidelines. You set or modify the ranges for either random method in the *Random Frame Sizes (with CRC)* dialog box.

Figure 8-13 illustrates this dialog box that is accessed by clicking the *Set Default* button in the *Custom Frame Sizes Per Flow (bytes with CRC)* dialog box.

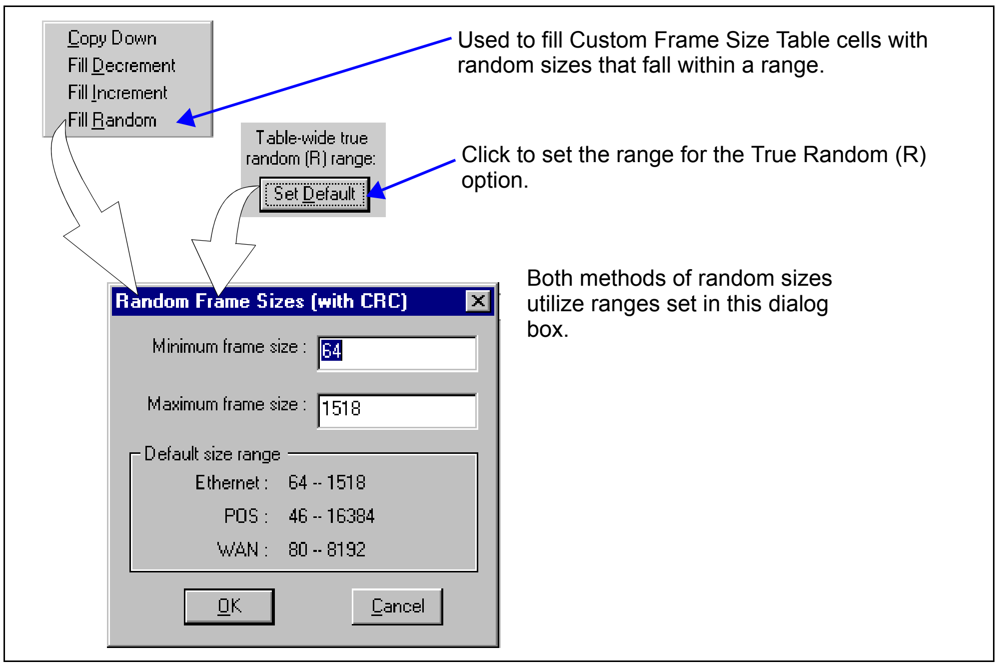

Figure 8-13. Random Frame Sizes

For a list of illegal frame sizes that SmartFlow allows, see *Appendix A, "Frame Size Limits"*.

For more information about how to use true random (R) frame sizes, including details about how to set the range for random sizes, refer to the online Help.

Global vs. Individual Cell True Random (R) Frame Size Ranges

The *Custom Frame Sizes Table* button allows you to access the cells set to R (True Random) that can use frame sizes from either:

- *A table-wide range* based on a specified range of legal sizes for that type of port.

  By clicking the *Set Default* button to access the *Random Frame Sizes (with CRC)* dialog box, the default frame size range and technology is set globally for every R cell in a table. This range applies to every *new and existing* R cell that uses table-wide ranges. Once you change the default, new and existing R cells that use the default now use the new default.

- *A range set for just the individual cell* based on the same technology and legal sizes.

  By double-clicking on an individual R cell to access the *Random Frame Sizes (with CRC)* dialog box, that cell can use a *different* size range from the default. This setting does *not* change the global defaults and is not affected by them if they change. (It must still use range allowed by SmartFlow.) *Figure 8-14* provides a comparison of these two ranges.

Note: If you change the random frame size range on a table-wide (global) basis, any existing sequences that also use global random settings that were based on a range for a *different* port technology will also change.

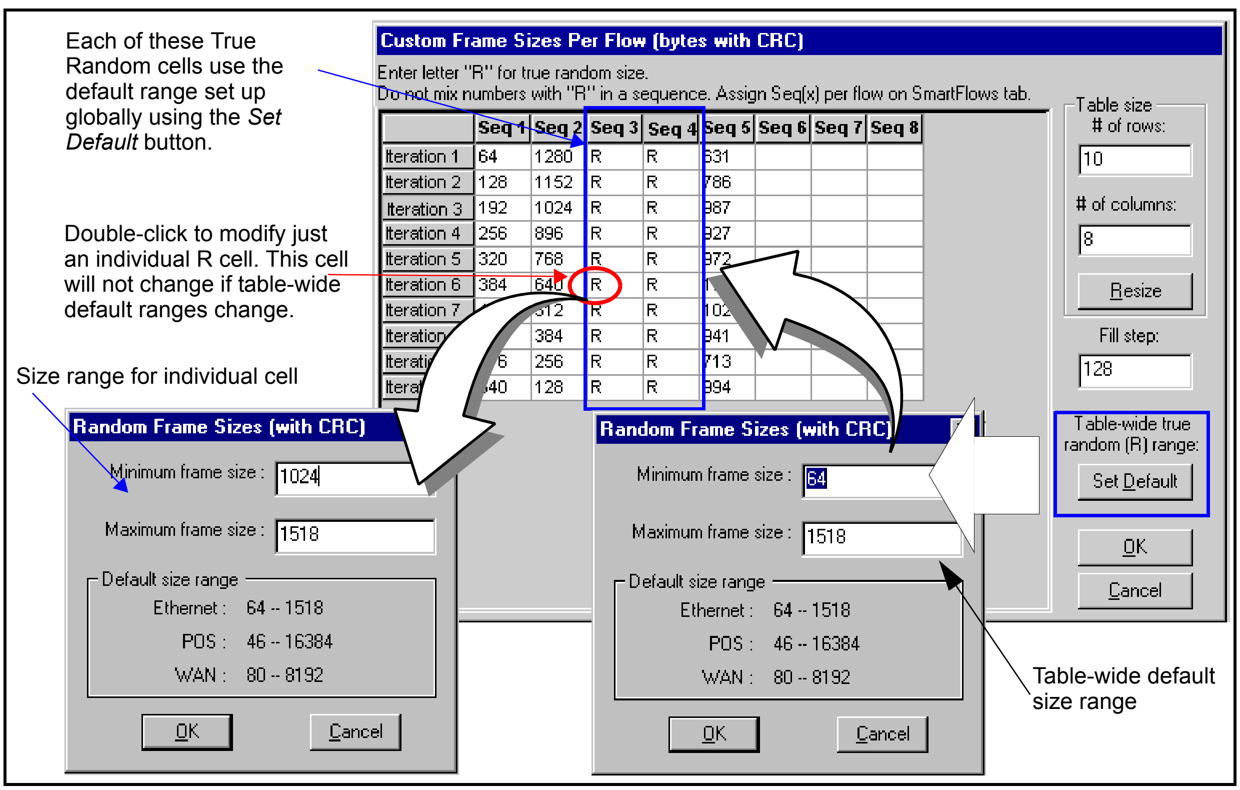

Figure 8-14. True Random (R) Table-wide vs. Individual Cell Default Frame Size Ranges

How Custom Frame Sizes Affects Test Results

When the *Custom frame sizes (per flow)* option is enabled on the *Test Setup>Test Iterations* tab, results show each flow by iteration.

For any flows that use true random (R) frame sizes, the *FrameSize* field shows zero, since each frame of the flow will be a different size. Group totals also show zero for this column–regardless of whether or not True Random was used–because each flow may use a different frame size. *Figure 8-15* illustrates this scenario.

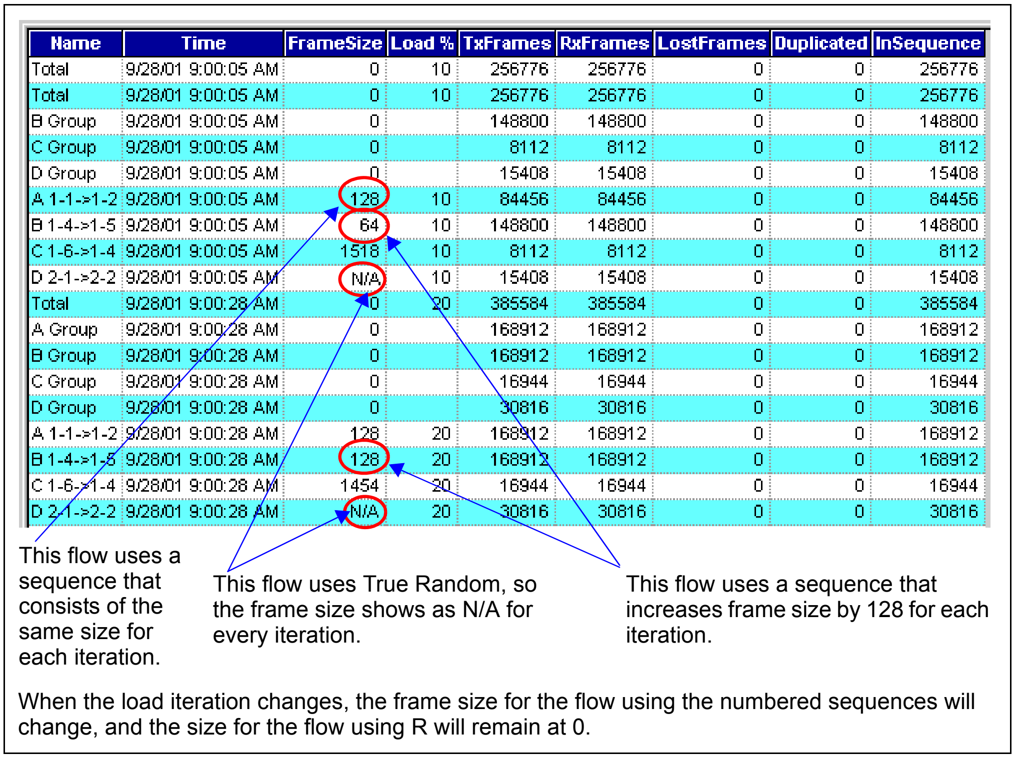

Name	Time	FrameSize	Load %	TxFrames	RxFrames	LostFrames	Duplicated	InSequence
Total	9/28/01 9:00:05 AM	0	10	256776	256776	0	0	256776
Total	9/28/01 9:00:05 AM	0	10	256776	256776	0	0	256776
B Group	9/28/01 9:00:05 AM	0		148800	148800	0	0	148800
C Group	9/28/01 9:00:05 AM	0		8112	8112	0	0	8112
D Group	9/28/01 9:00:05 AM	0		15408	15408	0	0	15408
A 1-1->1-2	9/28/01 9:00:05 AM	128	10	84456	84456	0	0	84456
B 1-4->1-5	9/28/01 9:00:05 AM	64	10	148800	148800	0	0	148800
C 1-6->1-4	9/28/01 9:00:05 AM	1518	10	8112	8112	0	0	8112
D 2-1->2-2	9/28/01 9:00:05 AM	N/A	10	15408	15408	0	0	15408
Total	9/28/01 9:00:28 AM	0	20	385584	385584	0	0	385584
A Group	9/28/01 9:00:28 AM	0		168912	168912	0	0	168912
B Group	9/28/01 9:00:28 AM	0		168912	168912	0	0	168912
C Group	9/28/01 9:00:28 AM	0		16944	16944	0	0	16944
D Group	9/28/01 9:00:28 AM	0		30816	30816	0	0	30816
A 1-1->1-2	9/28/01 9:00:28 AM	128	20	168912	168912	0	0	168912
B 1-4->1-5	9/28/01 9:00:28 AM	128	20	168912	168912	0	0	168912
C 1-6->1-4	9/28/01 9:00:28 AM	1454	20	16944	16944	0	0	16944
D 2-1->2-2	9/28/01 9:00:28 AM	N/A	20	30816	30816	0	0	30816

Figure 8-15. How Custom Frame Sizes Affect Test Results - Frame Loss Example

The Latency SnapShot test (*Figure 8-16*) results provide a list of each frame size at each load iteration:

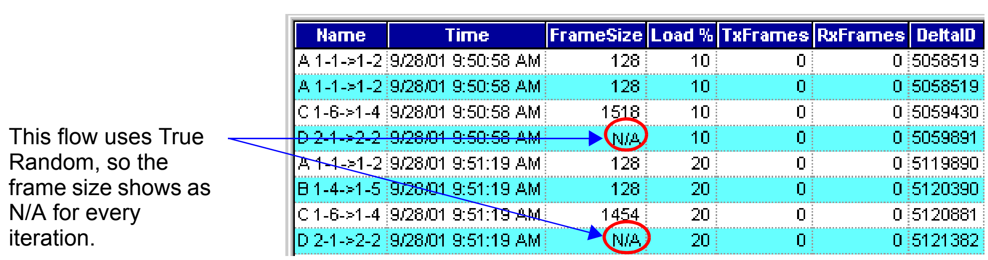

Name	Time	FrameSize	Load %	TxFrames	RxFrames	DeltaID
A 1-1->1-2	9/28/01 9:50:58 AM	128	10	0	0	5058519
A 1-1->1-2	9/28/01 9:50:58 AM	128	10	0	0	5058519
C 1-6->1-4	9/28/01 9:50:58 AM	1518	10	0	0	5059430
D 2-1->2-2	9/28/01 9:50:58 AM	N/A	10	0	0	5059891
A 1-1->1-2	9/28/01 9:51:19 AM	128	20	0	0	5119890
B 1-4->1-5	9/28/01 9:51:19 AM	128	20	0	0	5120390
C 1-6->1-4	9/28/01 9:51:19 AM	1454	20	0	0	5120881
D 2-1->2-2	9/28/01 9:51:19 AM	N/A	20	0	0	5121382

Figure 8-16. How Custom Frame Sizes Affect Test Results - Latency SnapShot Example

Flows using a True Random sequence display **N/A** for the frame size.

Using Frame Size Automation (Global for All Flows)

The *Frame size automation (all flows, with CRC)* option on the *Test Setup>Test Iterations* tab is a global way of varying the size of packets sent out during a test, in order to make the traffic more realistic. This method differs from the *Custom frame sizes (per flow)* option because *Frame size automation (all flows, with CRC)* uses the *same* frame size for every flow in the iteration.

The *Frame size automation (all flows, with CRC)* option allows you to vary the frame size either of two ways:

- *Step* option, which increases frame size in stepped amounts.
- *Custom* option, which varies frame size according to a list of sizes that may increase or decrease. See *"Customizing the List for Frame Size Automation (All Flows)" on page 239*.

You specify the frame size automation option at the *Test Setup Test Iterations* tab. For more information, see *"Test Iterations Tab" on page 217*.

Important: Using frame size automation will overwrite current frame sizes specified for flows. However, the original frame sizes are retained on the *Traffic* tab. To run a test at the original sizes, verify that the **Frame Size Automation (all flows, w/CRC)** button is not selected before you run the test.

If you set up flows using the Group Wizard, by using frame size automation you can use frame size automation together with the Group Wizard.

You set up frame size automation at the *Test Setup* > *Test Iterations* tab (*Figure 8-17*).

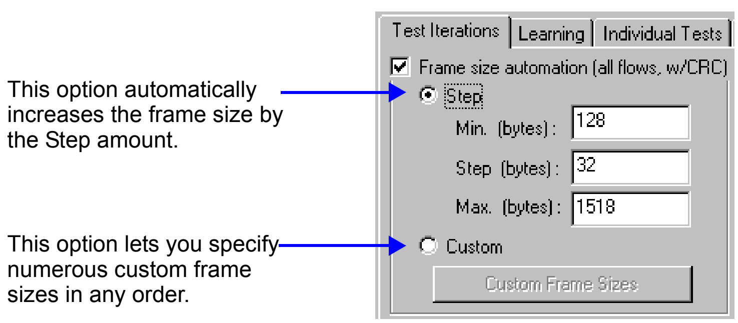

Figure 8-17. Frame Size Automation Option on the Test Iterations Tab

If the *Step (all ports)* option does not include the manner of varying frame sizes that you want, select the **Custom** option. See the section, *"Customizing the List for Frame Size Automation (All Flows)" on page 239* for more information.

To manually change frame sizes of multiple flows, see *"Changing Frame Sizes and Limiting Frame Length for Multiple Flows" on page 179*.

How Frame Size Automation Works

The general concept of how frame size automation works is as follows: SmartFlow transmits the same size of frames for each attempted traffic load and does not increment or change the frame size until the final load for the current frame size is complete.

Example:
Let us say you run the Latency test with three frame sizes: 128, 256, and 512. Assume the traffic load starts at 20%, steps up by 20%, and the maximum load is 80%. The test runs as shown here:

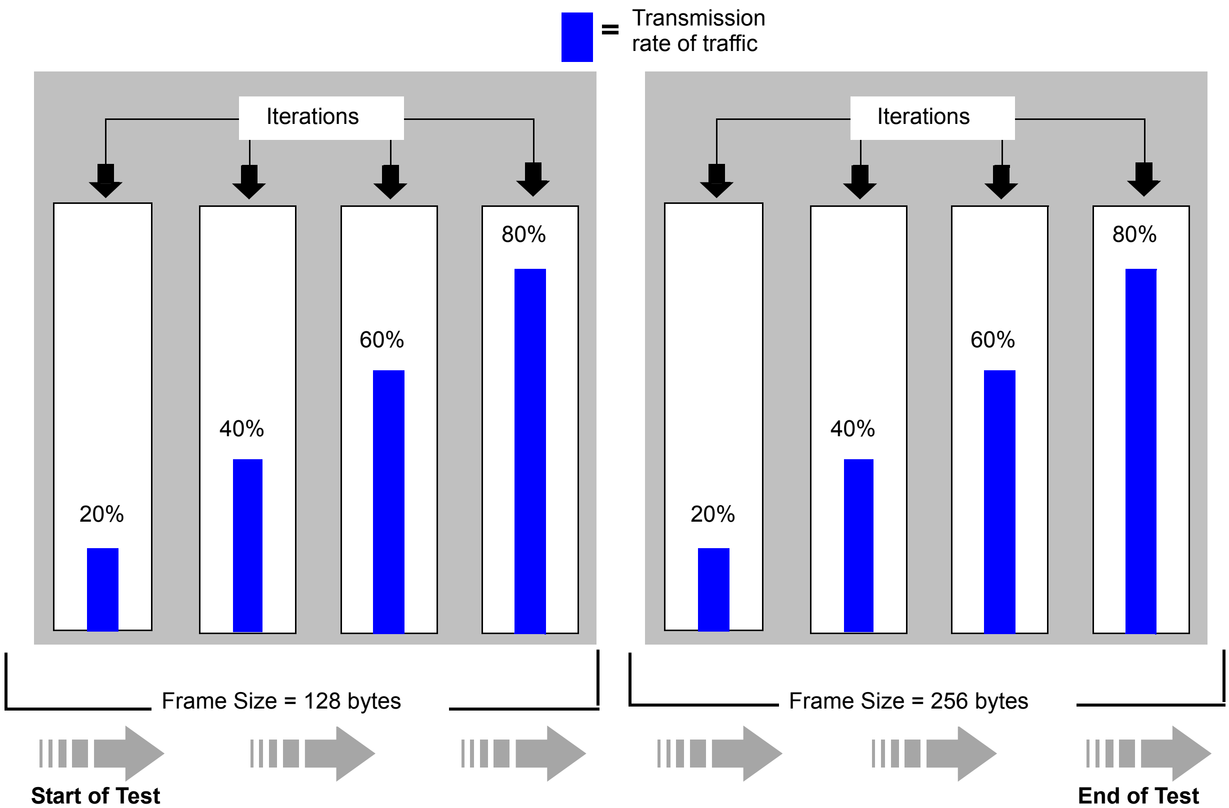

For the *Step* method of Frame Size Automation, SmartFlow does the following:

1. Transmits frames that are the size specified in the *Min.* field on the *Test Iterations* tab, at the minimum (initial load) for the test.
2. Continues to transmit the same size frames but increases the load until the maximum load is reached (for that frame size).
3. Increases the frame size by the number of bytes in the *Step* field on the *Test Iterations* tab.
4. Repeats steps 2 and 3 until the maximum frame size specified in the *Max.* field is reached. SmartFlow transmits frames at this maximum size for each load, but does not continue beyond this point.

For the *Custom* method of Frame Size Automation, SmartFlow does the following:

1. Transmits frames that are the first size at the top of the list in the Custom Frame Sizes dialog box, at the minimum (initial load) for the test.
2. Continues to transmit the same size frames but increases the load until the maximum load is reached (for that frame size).
3. Changes the frame size to the next size (it may be larger or smaller) in the list of custom frame sizes.
4. Repeats steps 2 and 3 for each size in the custom frame size list. SmartFlow transmits frames at this last size for each load, and then stops.

How Frame Size Automation Affects Test Results

When frame size automation is enabled, results are grouped first by *frame size* and then load for each flow. The more frame sizes that you use, the larger the results will be.

In *Figure 8-18*, notice that once the load reached 100% *Total* at the 128-byte *Frame Size*, the frame size increased to 256 bytes and started again at the lowest load, 10%.

The highest load for the 128-byte size.

The frame size increases to 256 bytes and the load starts at the beginning again.

Name	Frame Size	Load (%)	Throughput	Sent	Received	Lost	Loss (%)
Total		100.00000	100	844,584	844,584	0	0.00000
A TOS0		N/A	100	211,146	211,146	0	0.00000
A TOS2		N/A	100	211,146	211,146	0	0.00000
A TOS4		N/A	100	211,146	211,146	0	0.00000
A TOS6		N/A	100	211,146	211,146	0	0.00000
A T0 1-1->2-1	128	12.50000	100	105,573	105,573	0	0.00000
A T0 1-1->2-2	128	12.50000	100	105,573	105,573	0	0.00000
A T2 1-1->2-1	128	12.50000	100	105,573	105,573	0	0.00000
A T2 1-1->2-2	128	12.50000	100	105,573	105,573	0	0.00000
A T4 1-1->2-1	128	12.50000	100	105,573	105,573	0	0.00000
A T4 1-1->2-2	128	12.50000	100	105,573	105,573	0	0.00000
A T6 1-1->2-1	128	12.50000	100	105,573	105,573	0	0.00000
A T6 1-1->2-2	128	12.50000	100	105,573	105,573	0	0.00000
Total		10.00000	10	45,288	45,288	0	0.00000
A TOS0		N/A	10	11,322	11,322	0	0.00000
A TOS2		N/A	10	11,322	11,322	0	0.00000
A TOS4		N/A	10	11,322	11,322	0	0.00000
A TOS6		N/A	10	11,322	11,322	0	0.00000
A T0 1-1->2-1	256	1.25000	10	5,661	5,661	0	0.00000
A T0 1-1->2-2	256	1.25000	10	5,661	5,661	0	0.00000
A T2 1-1->2-1	256	1.25000	10	5,661	5,661	0	0.00000
A T2 1-1->2-2	256	1.25000	10	5,661	5,661	0	0.00000
A T4 1-1->2-1	256	1.25000	10	5,661	5,661	0	0.00000
A T4 1-1->2-2	256	1.25000	10	5,661	5,661	0	0.00000
A T6 1-1->2-1	256	1.25000	10	5,661	5,661	0	0.00000
A T6 1-1->2-2	256	1.25000	10	5,661	5,661	0	0.00000

Figure 8-18. How Test Results Appear with Frame Size Automation

The Latency SnapShot (*Figure 8-19*) test results provide a view for each rate of each frame size:

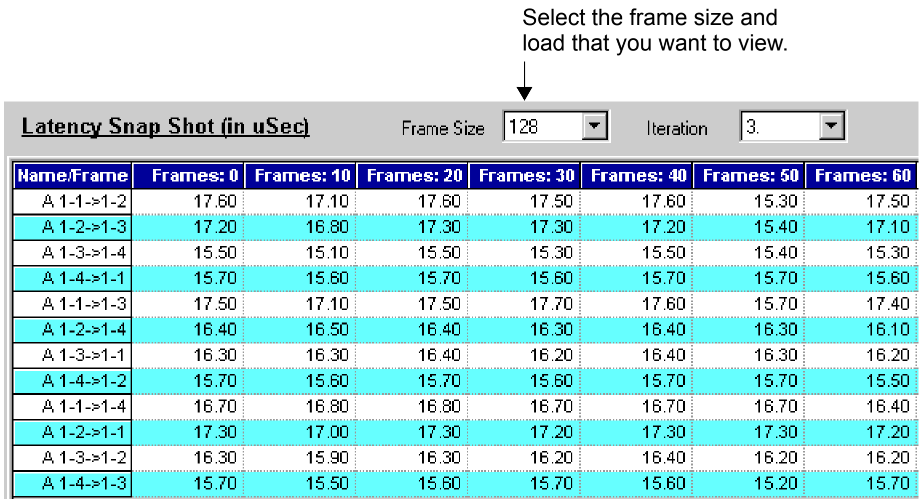

Name/Frame	Frames: 0	Frames: 10	Frames: 20	Frames: 30	Frames: 40	Frames: 50	Frames: 60
A 1-1->1-2	17.60	17.10	17.60	17.50	17.60	15.30	17.50
A 1-2->1-3	17.20	16.80	17.30	17.30	17.20	15.40	17.10
A 1-3->1-4	15.50	15.10	15.50	15.30	15.50	15.40	15.30
A 1-4->1-1	15.70	15.60	15.70	15.60	15.70	15.70	15.60
A 1-1->1-3	17.50	17.10	17.50	17.70	17.60	15.70	17.40
A 1-2->1-4	16.40	16.50	16.40	16.30	16.40	16.30	16.10
A 1-3->1-1	16.30	16.30	16.40	16.20	16.40	16.30	16.20
A 1-4->1-2	15.70	15.60	15.70	15.60	15.70	15.70	15.50
A 1-1->1-4	16.70	16.80	16.80	16.70	16.70	16.70	16.40
A 1-2->1-1	17.30	17.00	17.30	17.20	17.30	17.30	17.20
A 1-3->1-2	16.30	15.90	16.30	16.20	16.40	16.20	16.20
A 1-4->1-3	15.70	15.50	15.60	15.70	15.60	15.20	15.70

Figure 8-19. Latency SnapShot

Customizing the List for Frame Size Automation (All Flows)

Use the *Custom (all flows, with CRC)* option in the *Iterating across frame sizes* pane on the *Test Setup>Test Iterations* tab if you want to test any of these configurations:

- A customized list of frame sizes
- A specific order in which the frames are sent (such as 128, 512, 64)
- Descending frame order
- Numerous frame sizes.

You specify custom frame sizes on the *Custom Frame Sizes (all ports)* dialog box accessed by clicking the *Custom Frame Sizes List* button.

Figure 8-20 on page 240 shows the *Custom Frame Sizes (all ports)* dialog box.

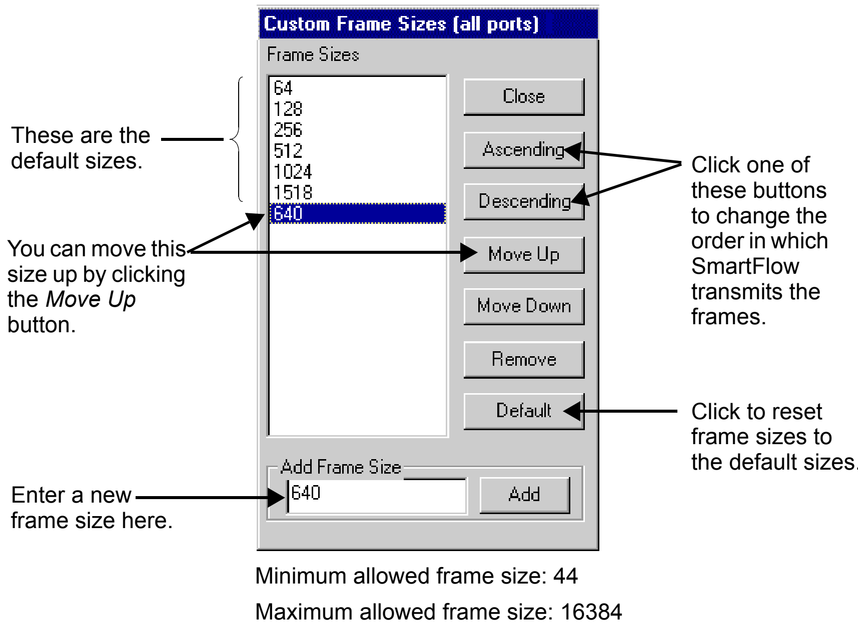

Figure 8-20. Custom Frame Sizes (all ports) Dialog Box

The default frame sizes that are provided span a broad range of size in ascending order. (The default sizes listed are arbitrary). To accept the default frame sizes, click the **Close** button.

SmartFlow uses frame sizes in the list from top to bottom. For example, let's say the list contains these sizes in the following order: 512, 128, and 256. SmartFlow first transmits 512-byte frames at all specified loads, then 128-byte frames at all specified loads, and then 256-byte frames. If you want the largest sized frame transmitted first, click the **Descending** button.

For more information about the *Custom Frame Size (all ports)* dialog box, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Using an IMIX Set of Frame Sizes for Flows

IMIX stands for Internet mix. It is a deterministic way of simulating real network traffic according to frame size usage. Some studies indicate that Internet traffic consists of fixed percentages of different frame sizes. IMIX traffic contains a mixture of frame sizes in a ratio to each other that approximates the overall makeup of frame sizes observed in real Internet traffic.

For example, one study showed that the top four frame sizes sent over the Internet (in terms of IP packet length) are:

- 40 bytes (comprising 56% of all traffic)
- 1,500 bytes (23% of all traffic)
- 576 bytes (17% of all traffic)
- 52 bytes (5% of all traffic).

The total is more than 100% due to rounding.

Using IMIX traffic allows you to test your device under more realistic conditions. You can base your IMIX frame sizes on samples from organizations that measure the makeup of Internet traffic. Alternatively, you may want to simulate some other usage/transmission pattern of certain frame sizes relative to one another.

IMIX Editor

Use the IMIX Editor to create a set of frame sizes and the corresponding percentages or weight that each frame size should occupy within an IMIX of flows. The IMIX Editor shows each IMIX set and its contents. It allows you to add, delete, or copy down IMIX sets. For whichever IMIX is highlighted, the window displays the corresponding frame sizes and their percentages. The total of the percentages is displayed at the bottom of the list.

The total of percentages for frame sizes in each IMIX set must add up to 100%. (If percentages add up to 99.99999, SmartFlow treats this as 100%.) Each IMIX set can contain up to 50 frame size/percentage pairs (lines). However, the Editor does not allow you to add any more frame sizes once the existing frame size percentages total to 100%. If you have only three frame sizes/percentages in an IMIX set and each size has equal weight (33.33333), SmartFlow allows a total of 99.99999%.

Only one random (R) frame size value is allowed per IMIX set.

To access the IMIX Editor:

1. Click the **Groups** tab.
2. Click the **IMIX Editor** button.
3. Use the editor to set up an IMIX set as illustrated in *Figure 8-21 on page 242*.

Some common IMIX sets are:

64 – 58.33333%

594 – 33.33333%

1518 – 8.33333%

Another common set is one that consists of equal weights for each frame size:

512 – 33.33333%

1518 – 33.33333%

9000 – 33.33333%

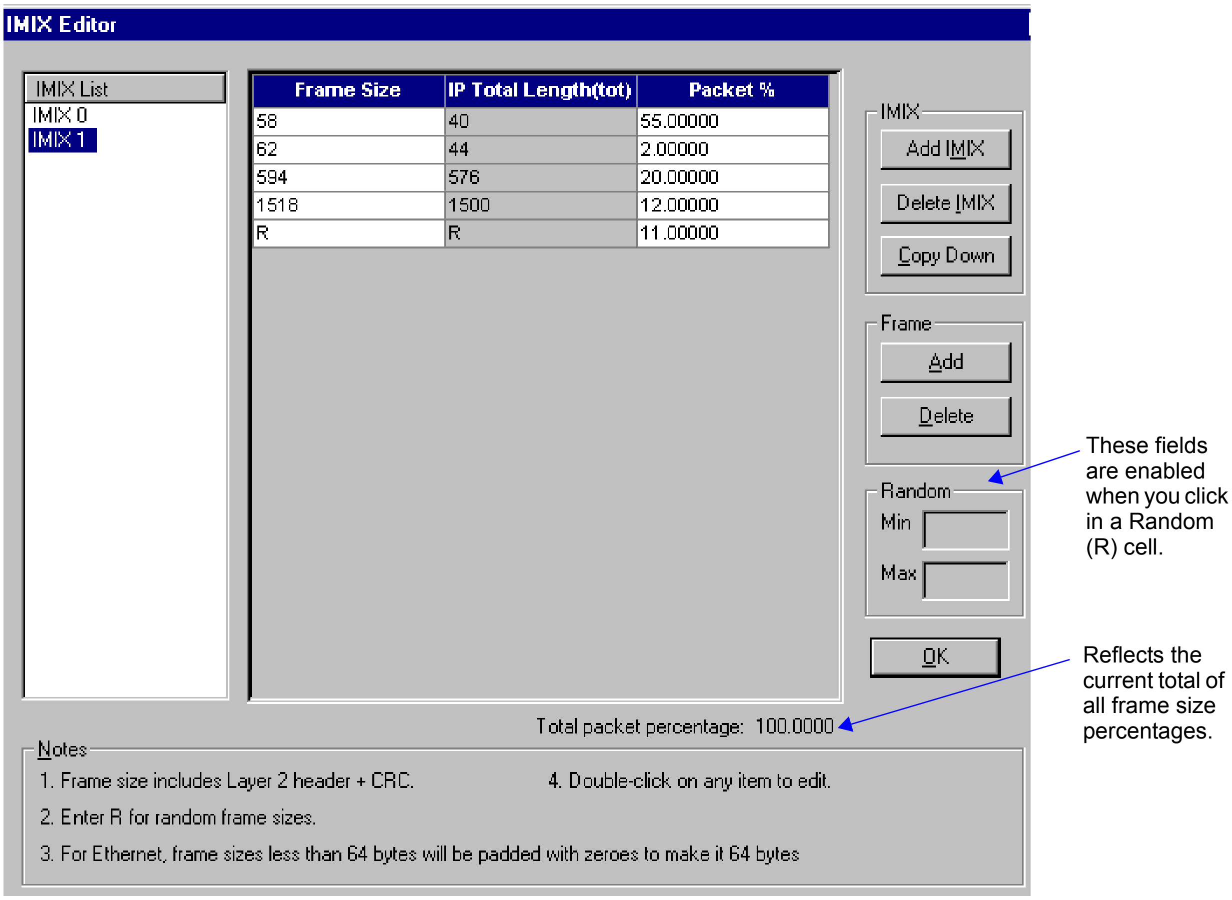

Figure 8-21. IMIX Editor

For field descriptions or detailed procedures, refer to the online Help.

Notes:

- Packet percentages are only for frame size ratios within that IMIX set.
- If you are creating IMIX flows with the Group Wizard, you must first set up an IMIX set, since one flow is created for each frame size in the set.
- It is not recommended that you manually set up IMIX flows from the *SmartFlow* tab. Unless your IMIX set contains a few frame sizes, manually setting up IMIX flows can be slow and confusing since you will have to set up one flow for each frame size. Instead, use the Group Wizard to create a set of IMIX flows.

How IMIX Works

In SmartFlow, an IMIX set (pattern) consists of a mixture of packet sizes and the percentages each packet size should share of the total traffic for that IMIX. For example, if Flow 1 - Frame Size 2 (256 bytes) of IMIX 1 is 55%, then 55% of all of the packets sent by all flows belonging to IMIX 1 will be 256 bytes. Think of each IMIX set as a separate entity. The packet size percentages in each IMIX set must total 100%.

SmartFlow automatically converts the frame size percentages that you specify in the IMIX set into the proper traffic loads so that the frame size ratio is maintained, regardless of the other flows on the port or the test load at a particular iteration. For example, if Flow 1 has a frame size of 64 bytes and the percentage in the IMIX is 35%, then 35% of the all frame sizes transmitted for that IMIX set will be 64 bytes, regardless of the load at any iteration.

Note: IMIX is supported on the TeraMetrics-based modules.

IMIX flows support cyclic flows, IP Multicast, VLAN tags, MPLS tags, and IPv6.

IMIX test is only supported on TeraMetrics-based LAN-3301A, LAN-3302A, LAN-3306A, LAN-3321A, LAN-3325A, LAN-3327A, POS-3505As, POS-3505AR, POS-3511As, POS-3519As, and POS-3519AR.

Note: If frames contain VLAN or MPLS tags, the minimum frame size to specify for one tag in the IMIX Editor is 56.

Creating IMIX Flow with the Group Wizard

When you create flows with the Group Wizard and use an IMIX set for the frame size, one flow is created for each frame size in a particular IMIX set. Since IMIX sets are tied to flows (one flow per IMIX frame size), you can assign multiple IMIX sets to multiple flows that are associated with the same transmitting port.

Example:

128	–	5%
256	–	10%
64	–	20%
1518	–	30%
1280	–	35%

The Group Wizard creates five flows, each of which transmits one of the five frame sizes in the IMIX. Frames at these sizes are transmitted in the ratio specified in the IMIX set.

Here are the sample results for the IMIX set with flows on Gigabit ports transmitting at 100% load.

Flow Name	Frame Size	Iload	Tx Frame Count	Rx Frame Count	Packet %
F1	128	0.7643	64,517	64,517	4.997231%
F2	256	2.85065	129,115	129,115	10.000736%
F3	64	1.73518	258,227	258,227	20.001239%
F4	1518	47.65544	387,340	387,340	30.001820%
F5	1280	46.99442	451,856	451,856	34.998974%

The **Packet %** (which is calculated from the other columns) is very close to the actual percentages specified in the IMIX set.

How Using an IMIX Set Affects Flow Rates

SmartFlow determines rates of IMIX flows on a per flow basis. Factors that affect the actual rates include:

- Whether or not there are any other non-IMIX flows on the transmitting port
- How the rates for these other non-IMIX flows are specified
- Rate of the test iteration
- Number of frames sent
- Frame size.

Each IMIX set contains a percentage for each frame size. If there are no other flows associated with the transmitting port, then the percentage is based on the total bandwidth available for that port in that test iteration.

If a transmitting port has only IMIX flows associated with it, each IMIX flow shares an even distribution of the port speed as its rate.

IMIX Flow Rates When Other Flows Exist with Other Rate Methods

If a transmitting port has multiple flows with different methods of rates specified (for example, rates per flow and even distribution), the remaining bandwidth for an IMIX set of flows is evenly distributed among all of the flows in each IMIX group as well as any individual flows set to even distribution. The flows using specific rates per flow take priority over the IMIX flows. If the rates of flows using specific rates per flow total 100% for the port, the IMIX flows have 0% rate.

Example: Here is a very simple example in which three IMIX sets each contain only one frame size. Assume that Port 1 is the same transmitting port for all of these flows:

Flow 1 using **IMIX 1**

Flow 2 using **IMIX 2**

Flow 3 using **IMIX 3**

These IMIX flow rates will be **evenly distributed**.

Flow 4 uses a **rate per flow of 10%.**

Flow 5 uses an **even distribution** for its rate.

Flow 6 uses a rate per flow of 15%.

The total percent of port load for Flows 4 and 6 equals 25%. Since the rates of IMIX flows are always evenly distributed across the remaining bandwidth, this allows Flows 1, 2, 3, and 5 to equally share 75% of the transmitting port (Port 1) bandwidth.

Manually Working with IMIX Flows

If you are using an IMIX set (weighted frame sizes) and plan on using many frame sizes, it is not recommended that you manually set up flows since it requires setting up one flow for each frame size, including an R (random) size. If you have many frame sizes in an IMIX set, use the Group Wizard to set up the IMIX flows.

Important: You can manually set up IMIX flows as long as you create one flow for every frame size in an IMIX set. If any frame size in an IMIX is missing a corresponding flow, the test does not run.

IMIX Behavior and Limitations

Using IMIX sets to create flows affects both the flow's frame size and its rate. The following behavior and limitations apply to IMIX flows:

- IMIX flows support cyclic flows, IP Multicast, VLAN tags, MPLS tags, and IPv6.
- If you select the *IMIX flow* checkbox on the *SmartFlows>Traffic* tab and any of the selected ports do not support IMIX, the flows associated with those ports are created as non-IMIX flows. You cannot use IMIX flows with frame size automation, custom frame sizes, or custom loads. If you enabled frame size automation, the presence of IMIX flows in the test will not prevent the test from running, but only the non-IMIX flows will change frame sizes. However, if you enabled custom frame sizes or custom loads, the presence of IMIX flows will stop the test from running.
- If you open a configuration file while connected to a chassis other than the original one, all IMIX flows become non-IMIX flows.

IMIX flows are similar to flows whose rate is set on a per flow basis in that the following limitations apply:

- On the *Groups* tab, you cannot have a combination of flows set to N/A together with flows set to even distribution or IMIX, and/or specific rates on the same port. Either all flows or no flows on the same port must be set to N/A.
- You cannot use the true random test loads together with IMIX flows. If the test includes this feature while IMIX flows exist, you will not be able to run the test. You must either disable the IMIX flows or not use any true random test loads.
- You can use frame size automation with rates per flow only when the rate unit is a percentage.

Random (R) frame size restrictions:

- Only one R (random) frame size is allowed per IMIX set. However, more than one IMIX set can be associated with a port.
- If multiple IMIX flows with an R frame size are on the same Tx port but the random minimum/maximum range differs for each one, SmartFlow prompts you to run the test with a readjusted random range that uses the lowest minimum and highest maximum frame size of the random ranges.
- You cannot run a Throughput test if any IMIX flows use an R frame size.

Creating IMIX Flows with the Group Wizard

Before you can use the Group Wizard to create an IMIX set of flows, you must set up an IMIX set of weighted frame sizes with the IMIX Editor.

To create an IMIX set of flows with the Group Wizard:

1. Select the **IMIX flow** checkbox on the **Group Wizard - Characteristics** page.
2. Choose an IMIX set.

The Group Wizard creates one flow for each frame size in a selected IMIX set.

Example:

	Flow Name	% / Custo
	☐ Flow 1	Even distri
	☐ Flow 2	Even distri
	☑ A 1->10-F1	IMIX
	☑ A 1->10-F2	IMIX
	☑ A 1->10-F3	IMIX
	☑ A 1->10-F4	IMIX

However, if you also select the **Generate multiple groups by** checkbox, the Group Wizard creates one flow for each frame size for each priority group.

Multiple transmitting ports can use the same IMIX sets. A transmitting port can have multiple IMIX sets associated with it.

Creating IMIX Flows with QoS or VLAN Priority

When you select the *Generate multiple groups by* checkbox on the *Group Wizard - Characteristics* page, the Group Wizard creates one flow for each frame size for each priority group. This results in multiple flows for each frame size. For example, if you selected three priority levels for DiffServ, SmartFlow generates three F1 (frame size 1) flows:

D0 - F1

D3 - F1

D7 - F1.

SmartFlow appends the frame size number (not size) from the IMIX set at the end of the flow name after the receiving port, such as **D0 2->12-F1**.

Example:

Let us say you chose to generate multiple groups by DiffServ values 0, 3, and 7. If you used an IMIX pattern with five frame sizes, the following flows may result for each priority group that you selected:

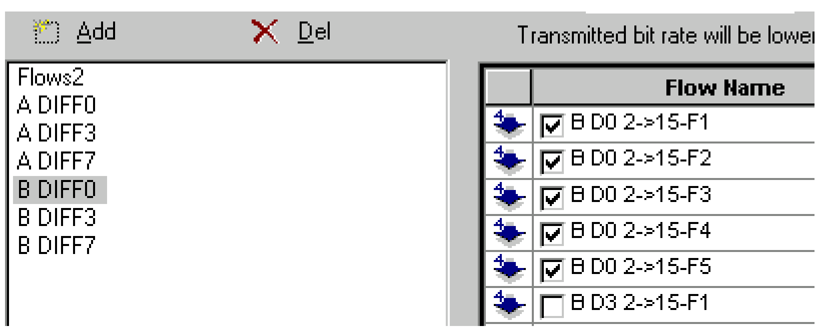

Since there are now multiple flows with the *same* frame size but each is in a different priority group, the total of each flow's packet size percentage equals the IMIX percentage for that frame size.

Example:

Assume that Flow 1 has a frame size of 64 bytes and the percentage in the IMIX for this frame size is 33%. The total of the frame size percentages transmitted for Flows D0 - F1, D3 - F1, and D7 - F1 together will equal 33% at 64 bytes of the all frame sizes that the port transmits.

Creating Multiple IMIX Flows with QoS or VLAN Priority

On the *Group Wizard - Multiple Flows* page, if you enter a value greater than *1* in the *Number of flows* field and also select the *Generate multiple groups by* checkbox on the *Group Wizard - Characteristics* page, SmartFlow creates multiple flows for each frame size in each priority group. For example, if you selected three priority levels for DiffServ and entered *2* in the *Number of flows* field, SmartFlow generates six F1 (frame size 1) flows:

D0 - F1-0

D0 - F1-1

D3 - F1-0

D3 - F1-1

D7 - F1-0

D7 - F1-1.

Creating Custom Frames

You can use the *Custom Frames Table* dialog box (*Figure 8-22 on page 251*) on the *Test Setup>Test Iterations* tab to define custom headers that you can subsequently assign to flows. Each custom frame you build can serve as a resource and can be reused by multiple flows.

Custom frame headers can contain values different from those in standard IP, TCP, UDP, and ICMP headers included in SmartFlow. You can define a custom Layer 2, Layer 3, or Layer 4 protocol header, which is often required to test the new protocols. Custom headers can be 64 or 128 bytes long.

You can also use the *Custom Frames Table* button to edit frame contents, so that there is only a Layer 2 frame with no IP header. This is useful when Layer 2 DUTs make forwarding decisions based on the IP addresses in the IP header. You may want to compare the DUT's treatment of Layer 3 frames to that of Layer 2 frames without IP addresses. You could do this by transmitting flows that contain a Layer 2 Ethernet header and as well as ones with no Layer 3 header.

Flows using custom frames have a special icon  beside them on the *SmartFlows* and *Groups* tabs to indicate that they are custom, since the flow may use IPv4 or IPv6.

SmartFlow provides a Layer 2 header template for the following types of frame headers:

- Raw Ethernet (no IP header)
- LLC SNAP
- POS.

If you use one of these templates, you can still change the Layer 2 header portion or add Layer 3 information.

You can also build a completely custom frame that is not based on any of the above templates.

By making flows that use custom frames cyclic flows, the flows can contain variable fields, thereby transmitting tens of thousands of unique values.

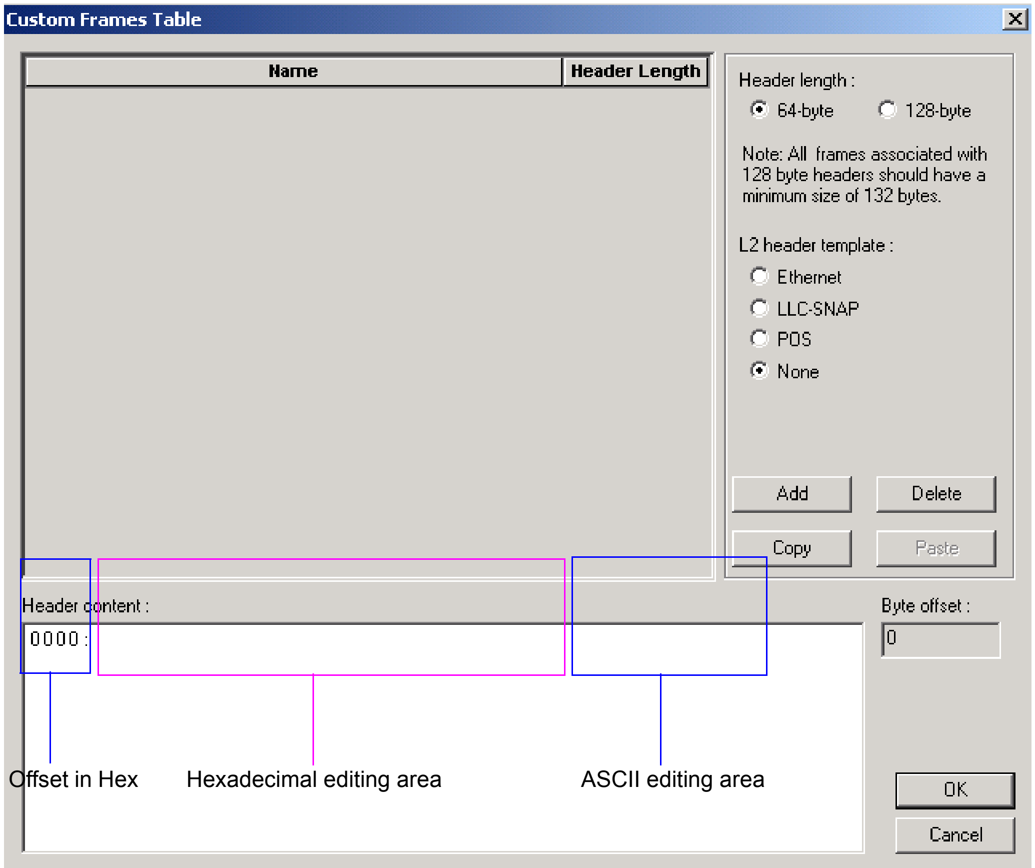

Figure 8-22. Use the Custom Frames Table Dialog Box to Create Custom Frame Headers

The *Custom Frames Table* dialog box allows you to construct different types of custom-built frame headers. You can add up to 2000 custom frame headers.

Tips:

- You can copy just a portion of the custom frame contents from one area to another within that header or to another frame header by using the *Ctrl* + *C* and *Ctrl* + *V* keys.

  You can copy the hexadecimal contents of a frame header from SmartWindow or any other application into the *Custom Frames Table* dialog box, as long as the byte format is xx xx xx.

For more information about custom frames, refer to the online Help. Press **F1** from the *Custom Frames Table* dialog box, or select **Help** from the menu bar.

Behavior and Limitations

- Card support: all SmartMetrics and TeraMetrics cards except for WAN.
- You can use custom frames with rates per flow, frame size automation, and custom frame sizes (but not random sizes).
- Cyclic flows are allowed on TeraMetrics-based transmitting ports.
- There is no Layer 3 ARPing or Neighbor Discovery for custom frames, but Layer 2 Ethernet learning is performed when an Ethernet port transmits custom frames (unless you specify no learning).
- Custom frames do not support data integrity checking (a port error).
- All Layer 3-related information on the *SmartFlows* subtabs (such as the *Traffic*, *VLAN* and *IP* tabs) is disabled for flows using custom frames. This is done to prevent information from being overwritten in the Layer 2 payload of the custom frame.

Varying Test Loads

Varying the test load (also known as *rate* or *percent of utilization*) is a way to make test traffic more realistic, since in real world traffic rates can vary greatly. For example, the rate at which a home user transmits traffic is often much slower than the rate at which a device on the Internet transmits traffic. (You can also vary the frame size to make traffic more realistic. See *"Varying Frame Sizes" on page 225*.)

You can vary the test load with each new test iteration in these ways:

- Per iteration only [*Step (all ports)* option]
  Specify a step (percentage) by which the load should increase with each test iteration. This increase affects all ports and all flows on each port. (To vary rates for individual flows per iteration, click the **Custom Loads Per Flow Table** button on the **Test Setup>Test Iterations** tab.) Click the **Step (all ports)** button. For more information, see *Table 8-3 on page 228*.
- Per port [*Custom (per port)* option]
  Create and assign a *custom sequence of loads* at which a port should transmit. Each load in the sequence is used in (corresponds to) one test iteration. Unlike the *per flow* method, this method attempts to evenly distribute the test load across all flows on the same port. Each flow equally shares the load.

  Custom loads can be a predefined or randomly generated percentage. You assign which sequence of sizes that you want a port to use for each test iteration at the *Test Load* field on the *Cards* tab. You can also specify a truly random (R) load percent for a cell or an entire sequence for a port. Click the **Custom (per port)** button on the **Test Setup>Test Iterations** tab to set up the sequences. For more information, see *"Custom Loads per Port" on page 255*.
- Per flow (*%/Rate unit/Custom* column)
  Specify for *each flow a percentage of total port load* or bit rate that the flow should utilize. Each flow can share either a different or evenly distributed proportion of the current test load for that port. Use rates per flow to transmit flows on a single port at different rates. Percentages remain constant across iterations. However, the rate per flow increases as a function of the increased overall test load.

  This option is only available for LAN-33xx, XLW-372x, and XFP-373x modules. You can also specify rates per flow by VLAN priority or class of service in VLANs. For more information, see *"Rates per Flow" on page 264*.

  Specify a rate for each flow at the *Groups* tab. Select the display by transmitting port, group, VLAN, and/or IP priority.

Notes:
- If you assign rates per flow, you cannot also use random frame sizes.
- You can assign rates per flow together with custom rates per port, as long as you do not use true random rates when defining custom rates. For more information about true random rates, see *"Using True Random (R) Test Loads" on page 259*.

- Per flow and per iteration [*Custom Loads (per flow)* option]

  Specify for *each iteration a percentage of total port load* that the flow should utilize. You can keep some flow rates constant across iterations while incrementing or decrementing other flow rates. This allows testing intra-port “throttling,” in which you can keep the rate of one flow or type of flow (such as voice traffic) constant at a given load across iterations, while other types of flows (such as those representing web traffic) increase or decrease.

  This option is only available for LAN-33xx, XLW-372x, and XFP-373x modules.

  Click the **Custom Loads per Flow Table** button on the **Test Iterations** tab or click in the appropriate cell in the **% /Rate unit/Custom** column on the **Groups** tab and select **Edit Sequences** to set up the sequences. For more information, see *“Rates per Flow and per Iteration” on page 275*.

Custom Loads per Port

When you vary rates *per port*, you assign which sequence of loads that you want a port to use for each test iteration in the *Test Load* field on the *Cards* tab. You can also specify a truly random ("R") load percent for a cell or an entire sequence for a port.

You can vary the load for a port in either of these ways:

- According to a predefined, *custom sequence* of rates
- *Randomly,* generating a rate for each test iteration.

Note: The *Custom* option for loads does not apply to a Throughput test since it uses a set number of test iterations, so the overall load cannot be determined.

How Custom Test Loads per Port Works

When you use custom test loads, each flow associated with a port is transmitted at the same rate (unless you also use rates per flow on the *Groups* tab). The rate changes when a new test iteration starts. (A *test iteration* is all flows in the test transmitting at a particular rate for the duration specified.) This is best understood with examples.

If you use custom test loads together with *rates per flow*, both the rate for each port and the rate for each flow on that port will vary for that test iteration. The rate for each port changes when a new test iteration starts, but the percentage that each flow utilizes of that rate will remain constant. For more information, see *"Rates per Flow" on page 264*.

If you use custom test loads together with c*ustom frame sizes*, both the rate and the frame size for each flow can change with each new test iteration. See *"Example 3: Custom Test Loads with Custom Frame Sizes per Flow" on page 257*.

You specify the sequences of loads for ports to use during a test at the *Custom Test Loads Per Port (% Utilization)* dialog box. See *"Specifying a Custom Test Load Sequence" on page 257* and *"Using True Random (R) Test Loads" on page 259* for more information.

Example 1: Custom Test Loads, Unidirectional

This example shows a section of how a test using custom test loads would run. Some of the test setup parameters include:
# of iterations used in test: 4
Seq. 1: 10%, 20%, 30%, 40%
Seq. 2: 80%, 85%, 90, 95%

***Port 1*:** Flows **A** and **B** use Seq. 1
***Port 2*:** Flows **C** and **D** use Seq. 2

SmartFlow transmits flows at the following rates as shown in *Table 8-4***.**

Table 8-4. Custom Test Loads - Unidirectional

Test Iteration	Port	Load	Load for Flow Associated with Port
1	Port 1	10%	Flow A - 10% Flow B - 10%
	Port 2	80%	Flow C - 80% Flow D - 80%
2	Port 1	20%	Flow A - 20% Flow B - 20%
	Port 2	85%	Flow C - 85% Flow D - 85%

The test continues changing rates at each new test load iteration, according to the predefined sequence of rates assigned to that port.

Example 2: Custom Test Loads, Bi-directional

Using the same test setup as Example 1, if you set up *bi-directional* flows, SmartFlow transmits flows at the rates shown in *Table 8-5*.

Table 8-5. Custom Test Loads - Bi-directional

Test Iteration	Port	Load	Load for Flow Associated with Port
1	Port 1	10%	Flow A - 10% Flow C- 10%
	Port 2	80%	Flow B- 80% Flow D - 80%
2	Port 1	20%	Flow A - 20% Flow C- 20%
	Port 2	85%	Flow B- 85% Flow D - 85%

When bi-directional is used, the ports still transmit flows at the same rate. In this example, Port 1 is now transmitting Flow A and C, instead of A and B. The test continues changing rates at each new test load iteration, according to the predefined sequence of rates assigned to that port.

Example 3: Custom Test Loads with Custom Frame Sizes per Flow

Some of the test setup parameters include:
# of iterations used in test: 4
Load Seq. 1: 10%, 20%, 30%, 40%
Load Seq. 2: 80%, 85%, 90, 95%

Frame Size Seq. 1: 64, 128, 512, 1028
Frame Size Seq. 2: 8192, 48, 764
Frame Size Seq. 3: 192, 192, 192, 192

Port 1: Transmits **Flow A,** which uses load Seq. **1** and frame size Seq. **1**
Transmits **Flow B,** which uses load Seq. **1** and frame size Seq. **2**
Port 2: Transmits **Flow C,** which uses load Seq. **2** and frame size Seq. **3**

If you use custom test loads together with custom frame sizes, SmartFlow transmits flows at the rates and frame sizes shown in *Table 8-6*.

Table 8-6. Custom Test Loads with Custom Frame Sizes per Flow

Test Iteration	Port	Load	Load/Frame Size for Flow Associated with Port
1	Port 1	10%	Flow A - 10% - 64 bytes Flow B - 10% - 8192 bytes
	Port 2	80%	Flow C - 80% - 192 bytes
2	Port 1	20%	Flow A - 20% -128 bytes Flow B - 20% - 48 bytes
	Port 2	85%	Flow C - 85% - 192

Notice how the ports still transmit all flows at the same rate, regardless of frame size.

For information about how Custom Frame Sizes works, see *"Using Custom Frame Sizes per Flow" on page 228*.

Specifying a Custom Test Load Sequence

You specify the custom (per port) test loads using the *Custom Loads Per Port Table* button on the *Test Setup>Test Iterations* tab. This dialog box allows you to specify a predefined series of load percentages to use at each iteration of a test, known as a *sequence*. You can assign a sequence to one or more ports. You define these sequences on the *Custom Test Loads Per Port (% Utilization)* dialog box, as shown in *Figure 8-23*.

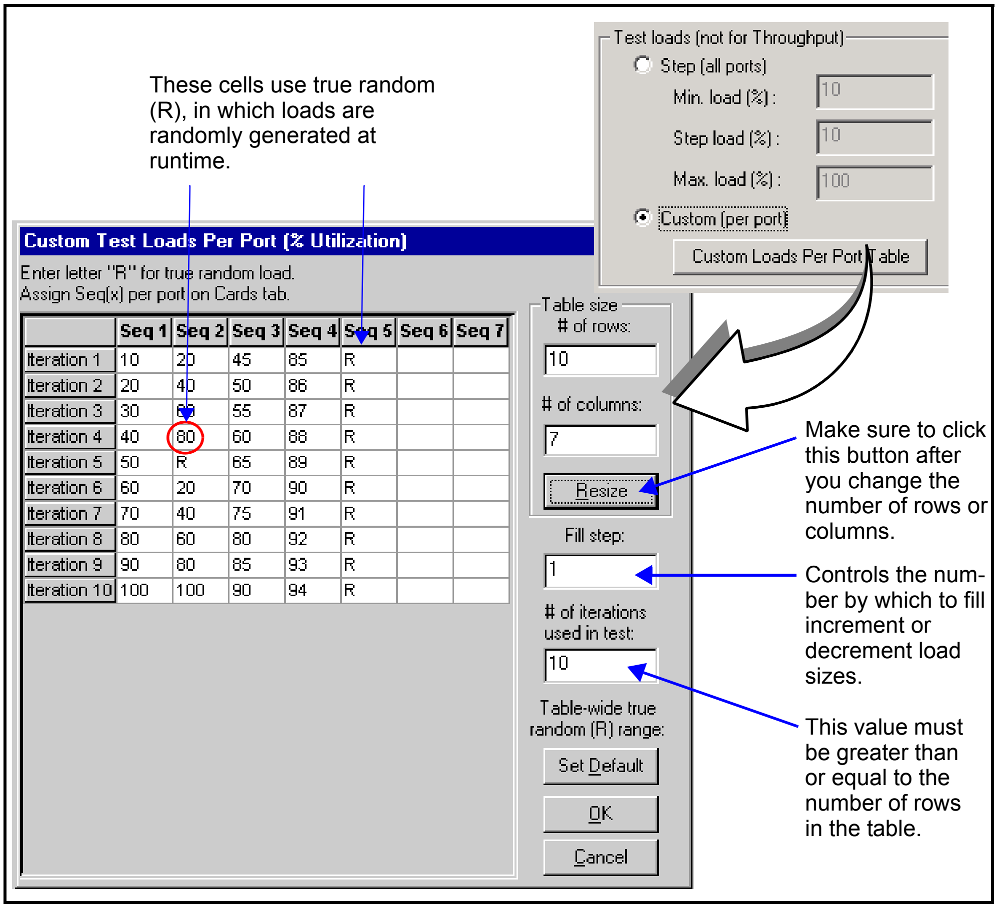

Figure 8-23. Using the Custom Test Loads Per Port (% Utilization) Dialog Box

Notes:
- Once you specify custom test load sequences, make sure that you also specify the correct load sequence for each port to use on the *Cards* tab.
- The value in the *# of iterations used in test* field must be greater than or equal to the number of rows in the table, or the test will not run.

Blank sequences can exist. However, all sequences must be either all blank or all filled. You can set up as many sequences as you want, depending on your computer's memory resources.

Important: If you change the number of rows or columns in the table, make sure that you click the **Resize** button for the changes to take effect.

On the *Test Setup>Test Iterations* tab, the *Custom Test Loads Per Port (% Utilization)* dialog box (accessed by clicking the *Custom Loads Per Port Table* button) is very similar to the *Custom Frame Sizes Per Flow (bytes with CRC)* dialog box (accessed by clicking the *Custom Frame Sizes Table* button). Both dialog boxes contain the same fields, offer many of the same options, and exhibit much the same behavior. The only differences between these two tables are that the *Custom Test Loads Per Port (% Utilization)* dialog box:

- Allows you to mix numbered and R cells within the same sequence.
- Allows you to set the number of test iterations, which is not dependent on any other settings.
- Has no fill random option.

For information about each field, see *"Specifying a Custom Frame Size Sequence" on page 229*. For information about copy/fill options or saving/reusing the table, refer to the online Help.

Using True Random (R) Test Loads

Instead of using load percentages in the *Custom Test Load Per Port (% Utilization)* dialog box, you can enter an "R" in any of the cells for a truly random rate that is generated at runtime by the card/module firmware.

Rates are based on the interframe gap (IFG) of frames transmitted from the port. Random rates will appear as N/A in test results since the rate is not uniform during the transmission.

True Random test loads differ from specified loads in a sequence in these ways:

- Load varies *from frame to frame* within a *stream* for the specified duration of the test load iteration.
- The *Custom Test Loads Per Port (% Utilization)* dialog box table cell(s) displays an R instead of the rate because rates are generated at runtime.
- Card/module firmware generates random rates with no pattern to them.
- Rates do not show up in results—N/A shows instead.

The following cards/modules support the true random (R) option: LAN-33xx, POS-3504/3505, POS-3510/3511, POS-3518/3519, XLW-3720/3721, and XFP-3730/3721. When assigning rate sequences to ports on the *Cards* tab, sequences with Rs do not appear in the *Test Load* field list of sequences unless the card/module supports the true random (R) option.

Tip: If you want to reproduce the test using the same random loads, enter a value in the *Random seed* field on the *Options>General* tab. When you rerun the test, make sure this same value is in the field. For more information about this field, see *"Setting up Global Application Preferences" on page 288*.

In *Figure 8-24*, notice how each IFG varies in Flows B and C, which use true random (R) loads, even within the same iteration.

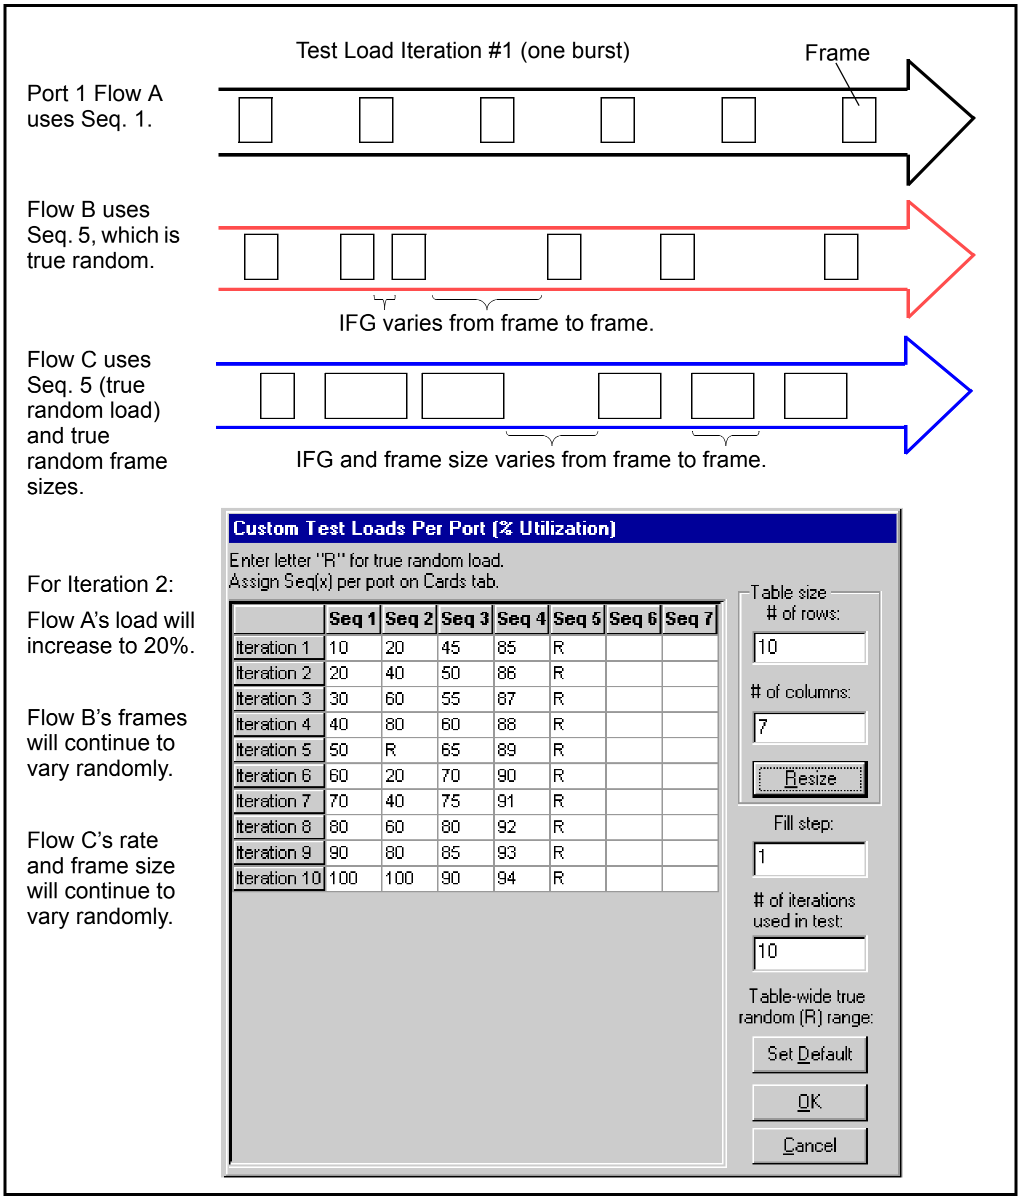

	Seq 1	Seq 2	Seq 3	Seq 4	Seq 5	Seq 6	Seq 7
Iteration 1	10	20	45	85	R		
Iteration 2	20	40	50	86	R		
Iteration 3	30	60	55	87	R		
Iteration 4	40	80	60	88	R		
Iteration 5	50	R	65	89	R		
Iteration 6	60	20	70	90	R		
Iteration 7	70	40	75	91	R		
Iteration 8	80	60	80	92	R		
Iteration 9	90	80	85	93	R		
Iteration 10	100	100	90	94	R		

Figure 8-24. Comparison of Custom and True Random Loads

For information about true random (R) for frame sizes, see *"Using Custom Frame Sizes per Flow" on page 228*.

Figure 8-25 illustrates how true random loads works from a port level.

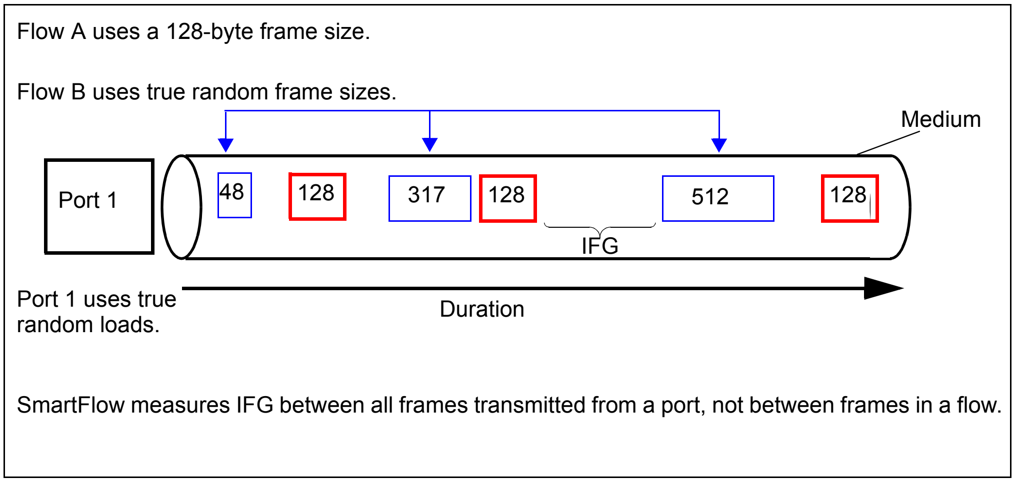

Figure 8-25. How True Random Load Works

Setting the Range for Random Rates

Random rates are generated based on the interframe gap. By setting a range of sizes for the interframe gap, you are in effect setting the range of random rates.

Use the following sample IFG size-to-rate equivalents as a guide to the IFG size that can be utilized.

IFG Size-to-Rate Equivalents

IFG	Load Equivalent (%)
96	100.000%
128	75.000%
192	50.000%
500	19.200%
1000	0.096%

For a procedure to set the range for true random (R) rates, refer to the online Help.

Global vs. Individual Cell True Random (R) Rate Ranges

Cells in the *Custom Test Loads Per Port (% Utilization)* dialog box set to true random (R) can use rates from either:

- *A table-wide range* based on a specified range of legal rates for that type of port.

  Click the **Set Default** button to access the *True Random (R) Rate Range* dialog box, the default frame size range and technology is set globally for every R cell in a table. This range applies to every *new and existing* R cell that uses table-wide ranges. Once you change the default, new and existing R cells that use the default will now use the new default.

- *A range set for just the individual cell* based on the same technology and legal rates.

  By double-clicking on an individual R cell to access the *True Random (R) Rate Range* dialog box, that cell can use a *different* rate range from the default. This setting does *not* change the global defaults and is not affected by them if they change. (It must still use the range allowed by SmartFlow.)

  A comparison of these two ranges is shown in *Figure 8-26*.

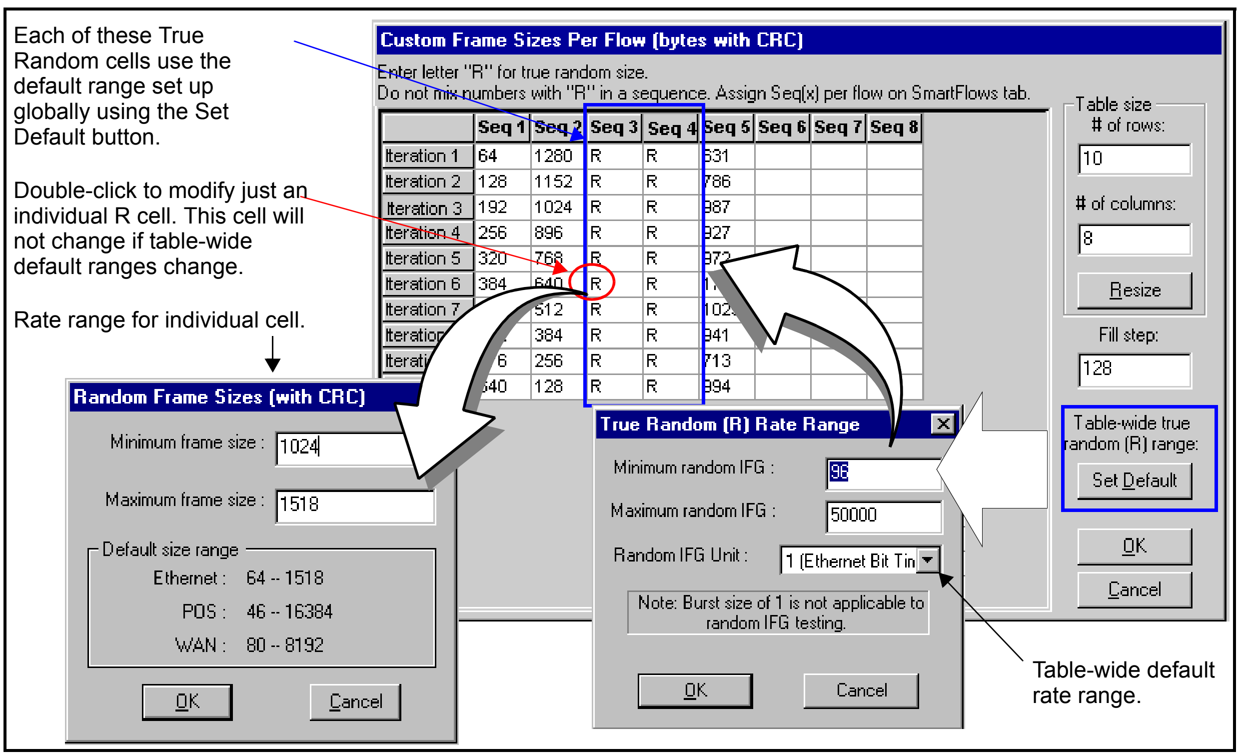

Figure 8-26. True Random (R) Table-wide Default Rate Ranges vs. Individual Cell Rate Ranges

For procedures to change the range of R cells, refer to the online Help.

How Custom Test Loads Affects Test Results

When the *Custom (per port)* test load option is enabled in the *Iterating across traffic load* pane on the *Test Setup>Test Iterations* tab, results show each flow by test load iteration. See *Figure 8-27*.

For any ports that use true random (R) loads, the *Load %* field shows N/A since each frame will have a different IFG and thus load. Groups totals show N/A for this column--regardless of whether or not true random was used--because each flow associated with a port may use a different rate.

With Custom Test Loads, group totals display N/A.

Frame Loss (%)

Name	Time	FrameSize	Load	TxFrames	RxFrames	LostFrames	Lost (%)	OutOfSequence	InSequence
A 1:17->1:16	11/12/01 3:30:12 PM	128	25.00	105564	0	105,564	100.00	0	0
Total	11/12/01 3:30:27 PM	128	N/A	506736	253368	253,368	50.00	0	253,368
A Group	11/12/01 3:30:27 PM	128	N/A	506736	253368	253,368	50.00	0	253,368
A 1:15->1:16	11/12/01 3:30:27 PM	128	30.00	126684	126684	0	0.00	0	126,684
A 1:15->1:18	11/12/01 3:30:27 PM	128	30.00	126684	0	126,684	100.00	0	0
A 1:17->1:18	11/12/01 3:30:27 PM	128	30.00	126684	126684	0	0.00	0	126,684
A 1:17->1:16	11/12/01 3:30:27 PM	128	30.00	126684	0	126,684	100.00	0	0
Total	11/12/01 3:30:42 PM	128	N/A	591216	295608	295,608	50.00	0	295,608
A Group	11/12/01 3:30:42 PM	128	N/A	591216	295608	295,608	50.00	0	295,608
A 1:15->1:16	11/12/01 3:30:42 PM	128	35.00	147804	147804	0	0.00	0	147,804
A 1:15->1:18	11/12/01 3:30:42 PM	128	35.00	147804	0	147,804	100.00	0	0
A 1:17->1:18	11/12/01 3:30:42 PM	128	35.00	147804	147804	0	0.00	0	147,804
A 1:17->1:16	11/12/01 3:30:42 PM	128	35.00	147804	0	147,804	100.00	0	0
Total	11/12/01 3:30:57 PM	256	N/A	45264	22632	22,632	50.00	0	22,632
A Group	11/12/01 3:30:57 PM	256	N/A	45264	22632	22,632	50.00	0	22,632
A 1:15->1:16	11/12/01 3:30:57 PM	256	5.00	11316	11316	0	0.00	0	11,316
A 1:15->1:18	11/12/01 3:30:57 PM	256	5.00	11316	0	11,316	100.00	0	0
A 1:17->1:18	11/12/01 3:30:57 PM	256	5.00	11316	11316	0	0.00	0	11,316
A 1:17->1:16	11/12/01 3:30:57 PM	256	5.00	11316	0	11,316	100.00	0	0
Total	11/12/01 3:31:12 PM	256	N/A	90576	45288	45,288	50.00	0	45,288
A Group	11/12/01 3:31:12 PM	256	N/A	90576	45288	45,288	50.00	0	45,288

Figure 8-27. Custom Test Loads in Test Results

Rates per Flow

When you vary the load (rate) per flow, you specify a percentage for the flow that is relative to the port's total utilization rate.

Percentages or bit rate units that you assign to flows (on the *Groups* tab) determine the proportion of port bandwidth for the flow in relation to other flows on that port. That proportion remains constant across iterations of the test. However, the actual rate that the flow utilizes changes as a result of the overall test load per port changing with each new iteration.

You can assign rates for flows on a port by:

- Specific percentages or bit rate units (such as Kbps), which are an uneven, user-defined distribution
- An even distribution of percentage amongst flows
- A combination of percentages/units and even distribution
- Internet mix (IMIX), which is an even distribution.

Percentages or bit rate units remain constant across iterations. However, the actual rate that the flow utilizes changes as a result of the overall test load per port changing with each new iteration. By contrast, when you do *not* assign rates per flow, each flow on the port sends roughly the same number of frames per second (based on the settings on the *Test Setup>Test Iterations* tab).

Notes:

- Rates per flow are only available for all LAN-33xx, XLW-372x, and XFP-373x modules.
- You cannot have a combination of flows set to N/A together with flows set to Even distrib. or IMIX, and/or specific rates on the same port. Either *all* flows or *no* flows on the same port must be set to N/A.
- You can use rates per flow with custom test loads and Custom frame size per flow. However, you cannot use rates per flow with true random test loads. You can use frame size automation with rates per flow only when the rate unit is a percentage.
- If the rate for a flow is less than 1 fps, test results may be invalid.
- You can use rates per flow with custom test loads, but not with true random test loads. You can use frame size automation and custom frame sizes per flow with rates per flow when the rate unit is a percentage, but not with true random frame sizes.

You assign rates per flow at the *Groups* tab once you have created some flows. For more information about how to specify rates per flow, see *"Rates per Flow When Specified in Bit Rate Units" on page 266*.

You can also assign rates to flows with the same VLAN, VLAN priority, and/or class of service.

Note: To disable the IMIX set for a particular flow, clear the *IMIX flow* checkbox on the *SmartFlows>Traffic* tab. In order to change the rate for an IMIX flow (for example, to a specific rate), you must delete the flow and then create it again without selecting IMIX for the frame size.

Tip: Once you finish assigning rates, if you want a less cluttered view of the *Groups* tab and no longer need to see the total rate for each port, change the flows *Sorted by* field to *Group*. The assigned rate total (excludes Even distrib.) lines are hidden.

How Rates per Flow Works

When you assign a rate per flow, each flow's percentage of the port load is the same for each test iteration. Although the flow's proportion of the port bandwidth stays the same, each time the test load increases with a new test iteration, the flow's rate also increases as a result.

When you assign a rate to a cyclic flow, that rate is evenly distributed among all of the cyclic variations. You cannot assign a rate for each variation of a cyclic flow, but the rate that you assign will apply to the base flow.

The rates per flow option is best understood by example. Assume the current test iteration is at 10% and you specify 5% for the *%/Rate unit/Custom* field on the *Groups* tab for a flow. This means the flow will use 5% of 10%, or .5% of the rate the port is capable of transmitting. If the next iteration is at 20%, the flow will use 5% of 20%, or 1% of the port's total utilization rate.

The concept of how rates per flow works is demonstrated more fully in *"Example: Two Transmitting Ports with Multiple Flows" on page 266*.

Example: Two Transmitting Ports with Multiple Flows

Assume that a test is set up with a minimum load of 10%. As the test progresses, SmartFlow transmits flows at the rates shown in *Table 8-7*.

Table 8-7. Two Transmitting Port with Multiple Flows

Test Iteration	Port	Load	%/ Rate unit/ Custom (Groups Tab)	Load for Flow Associated with Port
1	Port 1	10%	Flow A – 5% Flow B – 20% Flow C – Even distrib. Flow D – Even distrib.	Flow A – .5% Flow B – 2.0% Flow C – 3.75% Flow D – 3.75%
	Port 2	10%	Flow E – N/A Flow F – N/A	Flow E – 10% Flow F – 10%
2	Port 1	50%	Flow A – 5% Flow B – 20 Flow C – Even distrib. Flow D – Even distrib.	Flow A - 2.5% Flow B - 10% Flow C – 18.75 Flow D – 18.75
	Port 2	50%	Flow E – N/A Flow F – N/A	Flow E – 50% Flow F – 50%
3	Port 1	100%	Flow A – 5% Flow B – 20% Flow C – Even distrib. Flow D – Even distrib.	Flow A – 5% Flow B – 20% Flow C – 37.5. Flow D – 37.5.
	Port 2	100%	Flow E – N/A Flow F – N/A	Flow E – 100% Flow F – 100%

Rates per Flow When Specified in Bit Rate Units

When specifying rate per flow in bit rate units, the value that you specify on the *Groups* tab is based on 100% port load. For example, if you specify 800 fps for Port A, the 800 fps will be achieved when test load reaches 100%. Let us say the first iteration is 10% load, the second iteration is 20% load, and the third iteration is 30% load. The actual fps for the first iteration is 10% of 800=80 fps, the fps for the second iteration is 20% of 800=160 fps, and the fps for the third iteration is 30% of 800=240 fps.

The calculation for Mbps, Kbps and bps is based on:

Preamble + Layer 2 Header + Payload +CRC + IFG

The IFG is fixed at 96 bits.

Specifying Rates per Flow

Once you have created some flows, you can specify the rate per flow on the *Groups* tab. Before you enter rates, you may want to use the Flow Rate Calculator first to estimate the allowable percentage per flow on that port. You can then copy the cells from the appropriate bit rate unit column or the *Util(%)* column in the Flow Rate Calculator to the *%/Rate unit/Custom* column on the *Groups* tab. See *"Calculating Rates per Flow on the Flow Rate Calculator" on page 271* for more information.

The *%/Rate unit/Custom* field on the *Groups* tab represents that flow's proportion of the port's total available utilization rate during an iteration. Even distribution (Even distrib.) divides the available bandwidth that remains for the port after all specified rates are totaled and distributes it evenly across those flows. Any IMIX flows also share an even distribution with the port's remaining bandwidth.

To specify rates per flow by IP priority in VLANs, see *"Setting Rates per Flow by IP Priority in VLANs" on page 270*.

Prior to assigning a rate, SmartFlow displays N/A as the rate for each flow in the *%/Rate unit/Custom* column on the *Groups* tab, as shown in *Figure 8-28*.

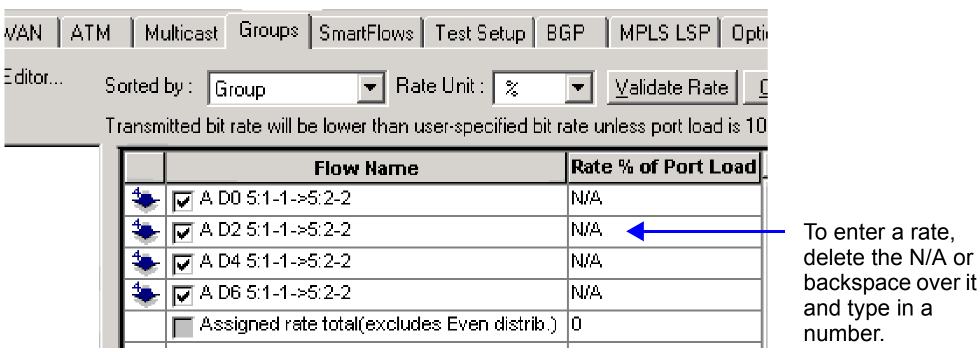

Figure 8-28. Groups Tab Prior to Assigning Rates per Flow

N/A indicates that the transmitting port for that flow transmits frames for each flow on that port in a round robin-fashion with no proportional allotment of the port load to any one individual flow. It is the default for new flows. Once you enter a percentage or "e" for even distribution for one of multiple flows on a port, SmartFlow automatically changes the remainder of the flows to even distribution. Any new flows that you add to that port will then, by default, also be set to even distribution.

After you make an addition or change, click the **Validate Rate** button for each port. A message will display that informs you whether or not the flows on this port exceed the bandwidth (oversubscribe). This also updates the *Assigned rate total (excludes even distrib. & sequences)* field to reflect the current total percent. All flows on a port are validated together.

If the message box indicates that port bandwidth is oversubscribed, lower the percentage or bit rate of one or more of the flows on the port.

Tips:

- Once you finish assigning rates, if you want a less cluttered view of the *Groups* tab and no longer need to see the total rate for each port, change the *Sorted by* field to **Group**. The assigned rate total (excludes even distrib. & sequences) lines will be hidden.
- If you want the flow's assigned rate percentage of port load to be an actual, fixed percentage (as opposed to relative percentage or proportion) of the port's total bandwidth, set the minimum load to 100%. For more information, refer to the online Help.

See *Figure 8-29* and *Figure 8-30 on page 269*.

Notes:

- You cannot have a combination of flows set to N/A with flows set to Even distrib. and/or specific rates on the same port. Either *all* flows or *no* flows on the same port must be set to N/A.
- You can use rates per flow with custom test loads. However, you cannot use rates per flow with frame size automation, custom frame size per flow, or true random test loads.
- If the rate for a flow is less than 1 fps, test results may be invalid.

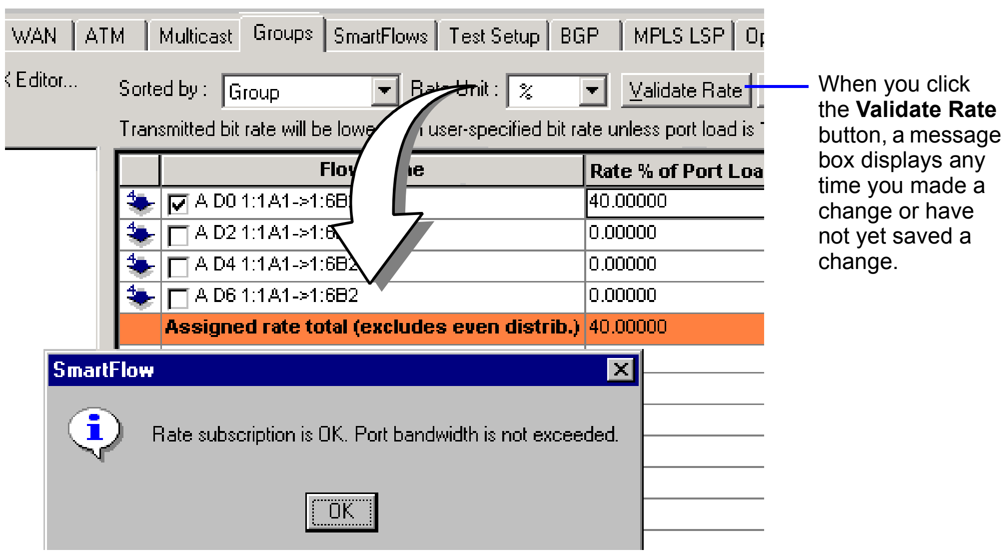

Figure 8-29. Groups Tab for One Group of Flows Prior After Assigning Rates per Flow

The *%/Rate unit/Custom* field represents that flow's proportion of the port's total available utilization rate during an iteration. Even distribution (even distrib.) divides the available bandwidth that remains for the port after all specified rates are totaled and distributes it evenly across those flows.

Figure 8-30 on page 269 shows grouped flows on the same port and non-grouped flows on different ports with assigned rates.

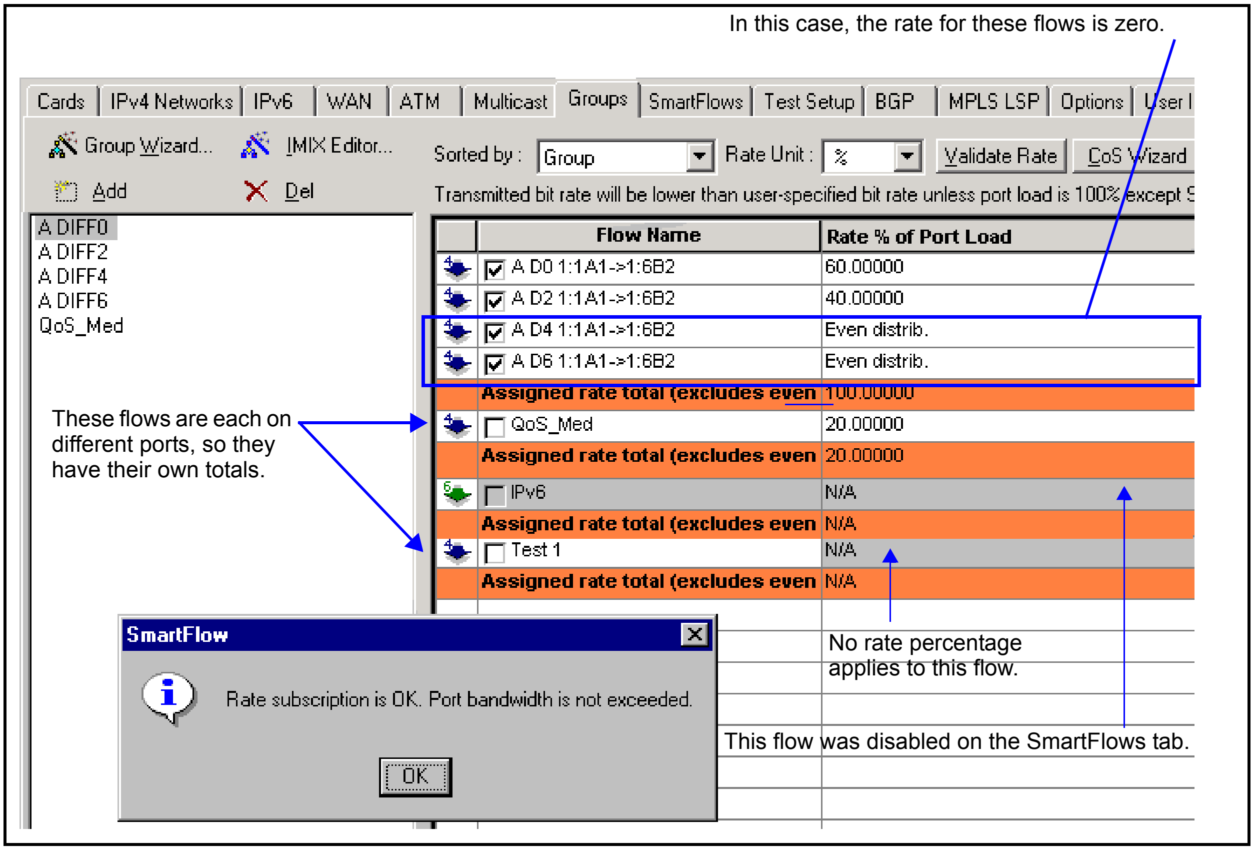

Figure 8-30. Groups Tab for Multiple Flows (Grouped and Non-grouped) After Assigning Rates

The total of the rates specified for each flow on a port is displayed on the *Assigned rate total (excludes Even distrib.)* line. Total assignment for a port cannot exceed 100%.

Flows that are grayed out are either disabled on the *SmartFlows* tab or the transmitting port is not a LAN-33xx, XLW-372x, or XFP-373x module.

Disabled flows with a specified rate have no effect on the port load. If the port has other flows that are still enabled, any of them set to *Even distrib.* now get a higher percentage of the port load since more bandwidth is available.

For information about how to copy, increment, or decrement rates for multiple flows, refer to the online Help.

Setting Rates per Flow by IP Priority in VLANs

If you provision bandwidth to provide various service levels to customers in your network, you can simulate this in SmartFlow by setting the rates of flows according to these same service levels. SmartFlow allows you to automatically set rates for IPv4 flows according to the flow's port/VLAN and IP priority (ToS or DiffServ), if the flow's transmitting port is in a VLAN. For example, you can set the rate for all flows with TOS 1 that is in VLAN 5 to 20%.

You can set rates per flow by port/VLAN and IP priority in two ways:

- Automatically and at the same time that flows are created, using the Group Wizard and the *Rate Setting for ToS/Diffserv* dialog box. This method also allows you to set rates per flow by VLAN according to subnet in addition to port.
- After flows are created, using the CoS Wizard on the *Groups* tab.

The CoS Wizard allows you to set rates for flows according to their class of service (CoS) number and VLAN. This saves a lot of time that would otherwise be needed to calculate and configure variations of common network scenarios based on VLANs and CoS.

(To simultaneously create flows and set rates per flow by VLAN and IP priority, use the Group Wizard. See *"Setting Rates per Flow by VLAN and IP Priority While Creating Flows" on page 153*.)

To access the CoS Wizard:

1 Click the **Groups** tab.
2 Click the **CoS Wizard** button.

First, you may want to sort and view the flows on the *Groups* tab by VLAN and/or IP priority (CoS).

For each flow with a transmitting port on a VLAN, use the CoS Wizard to assign (in order) the following:

- The percentage of port rate that VLANs on selected ports should utilize.
- The percentage that a ToS precedence or Diffserv class should utilize of each VLAN.

All flows in a particular VLAN that have been assigned a particular CoS number are transmitted at the specified rate. All flows with the same CoS in the same VLAN are transmitted at the same rate. The actual rate is incremented per the test load step or custom rate table, but the proportion of the rate per flow remains the same.

For example, if you specify:

- **VLAN 1 Rate = 10%** it means that all flows on VLAN 1 are transmitted at 10% of the current test load for the port.
- **ALL TOS 1 = 20%** it means that all flows in VLAN 1 with a TOS 1 priority are transmitted at 20% of 10%.

Once you use the CoS Wizard, the rates per flow appear on the *Groups* tab where you can sort and view them by VLAN and/or IP priority.

Notes:

- All of the CoS-based flow rates within a VLAN cannot add up to more than a 100, which represents 100% of the VLAN to which they belong. All of the VLAN rates cannot add up to more than 100, which represents 100% of the port rate.
- If you do not set the IP priority for a flow, the default is 0 priority.

For more information about the CoS Wizard, press **F1** from any of the CoS Wizard pages, or select **Help** from the menu bar.

Calculating Rates per Flow on the Flow Rate Calculator

Many DUTs are provisioned for QoS by setting the bandwidth in Mbps or Kbps. The Flow Rate Calculator allows you to convert these bit rate units into percentages or from percentages back to units that can be copied to the *Groups* tab. (There is no correlation between the values in the calculator and the *Groups* tab.) Use the calculator to calculate flow rates for *one port at a time*. Each flow listed in the table on the Flow Rate Calculator is generated by the same port.

Note: The contents of the Flow Rate Calculator can be saved, allowing you to reload the data from the file.

For Ethernet ports, use the Flow Rate Calculator to:

- Estimate the total allowable percentage for each flow transmitted by a port, at a specific test load and port speed.
- Convert between Kbps and Mbps.
- Select which *portions* of the packet to be recorded in Kbps or Mbps in order to identify overhead.

For ATM, use the Flow Rate Calculator to:

- Support ATM speeds of OC3c (155.52 Mbps) and OC12c (622.08 Mbps).
- Calculate the ATM speed from an equivalent frame rate for Ethernet.
- Produce a Kbps setting for use on the ATM card to allow identical fps values to be sent between ATM and Ethernet.
- Include/exclude parts of the ATM frame.
  The frame format is as follows:
  LLC Header + User Data (Frame Size in Bytes) + Padding (up to 47 Bytes) + AAL5 Trailer (8 Bytes)
- Convert between cells/second and frames/second. Conversions carry across the entire row, so given an input in a particular cell, equivalent values in other columns are calculated and displayed.

  Note: The user data is the frame size in bytes that is input into the table grid. The *User Data* and *Padding* fields are always enabled and cannot be toggled. The *Padding* field is always enabled in order to guarantee that the ATM frame is an integer multiple of 48 bytes since an ATM cell has 48 bytes of payload and a frame must fit in an integer number of cells.

- Show both the frame rate as well as the cell rate for ATM. Any change in the frame rate will adjust the cell rate accordingly and vice versa.
- Alter the *Frame Portion Selections* fields by either toggling the *AAL5 Trailer*, by selecting an encapsulation method for ATM (such as *LLC Routed IP*), or by selecting/clearing the *Sonet Overhead* checkbox.

  Note: The LLC encapsulation (routed/bridged) field values enable the *LLC Header* frame portion selection, and the VC encapsulation (routed/bridged) field values disable the *LLC Header* frame portion selection.

You can determine the rate for the flow by percent of utilization, frames per second, and bits per second. By specifying one unit, the calculator calculates the other two equivalent measurements.

Note: Bits per second and frames per second rates that display in the calculator are based on the current port load specified in the calculator, which is a fixed (static) rate. When the rest actually runs, the bit and frame rates change with each test load iteration.

Once the rates per flow for a port are calculated, copy the rates directly from any of the bit rate unit (*Mbps/Kbps* or *Frames/Sec*) or *Util(%)* cells in the calculator to the appropriate cells in the *%/Rate unit/Custom* column on the *Groups* tab. If you copy cells that are in bit rate units, that information is automatically converted to a percentage on the *Groups* tab. Then select the rate unit by which to view the flow rates.

For procedures on how to copy rates to the *Groups* tab, press **F1** at the Flow Rate Calculator, or select **Help** from the menu bar.

Constraints and Tips

Some constraints and tips relating to the Flow Rate Calculator are as follows:

- The frame size for Ethernet speed must be between 64 and 16388 bytes.
- The minimum frame size for ATM depends on the encapsulation type.
  - For LLC Routed IP, the minimum is 46 and the maximum is 32767.
  - For LLC Bridged IP, the minimum is 62 and the maximum is 32767.
  - For VC Muxed Routed, the minimum is 38 and the maximum is 32767.
  - For VC Muxed Bridged, the minimum is 54 and the maximum is 32767.
- Utilization must be between 0 and 100%.
- The *Kbps/Mbps*, *Frames/Sec*, and *Cells/Sec* columns must be between 0 and the maximum bps, pps, and fps, respectively, which depends on the frame size and speed selected.
- Click the **Calculate** button when a change is made in the Frame Portion Selections fields to enable/disable parts of ATM/Ethernet frame.

- To copy and paste the *Kbps* or *Mbps* column value into the *Speed* column on the *ATM* tab, press **Ctrl-C/Ctrl-V**.
- Any file that stores the contents of the calculator must not be modified outside the application. (The file only contains the selected column and column frame size.) The default extension is .cal.
- When loading the contents in the calculator from a file, the table size must be large enough to accommodate the contents in the calculator.

To access the Flow Rate Calculator:

1 Select **Tools>Flow Rate Calculator** from the menu bar.

Figure 8-31 on page 273 shows the Flow Rate Calculator.

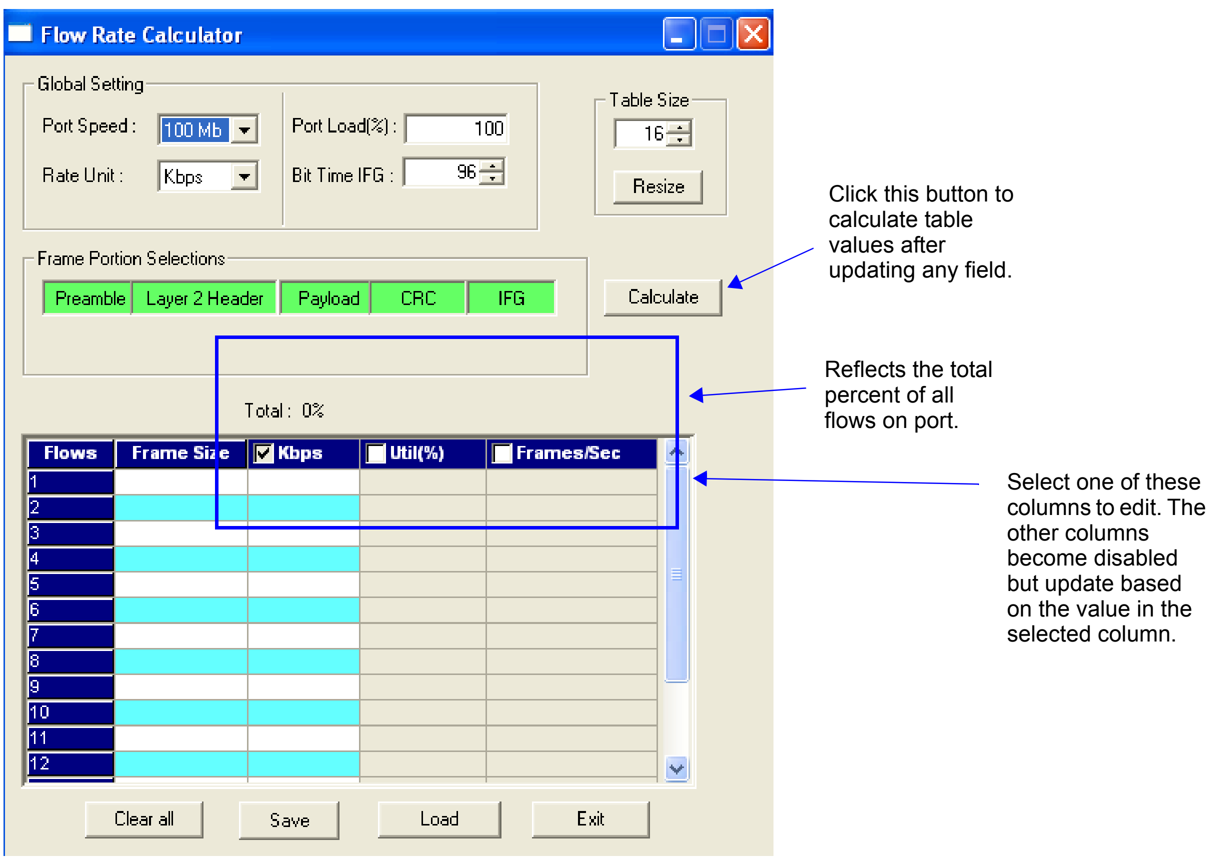

Figure 8-31. Flow Rate Calculator

For a description of each field and detailed procedures on how to use the Flow Rate Calculator, press **F1** at the Flow Rate Calculator, or select **Help** from the menu bar.

How Rates per Flow Affect Results

When a test is set up with rates assigned per flow, test results display each flow with a different rate within the same test iteration. *Figure 8-32* shows an example of Jumbo detailed test results with rates per flow. Notice how the rates vary in the *ILoad* column from flow to flow within the same iteration.

Jumbo (in uSec) Page 1 of 1

Name	Time	FrameSize	ILoad	TxFrames	RxFrames	LostFrames	StdDeviation	MinLatency	AveLat
Total	10/18/02 11:35:09	N/A	60.00000	956	956	0	62.00	23.30	N/A
A Group	10/18/02 11:35:09	N/A	N/A	956	956	0	62.00	23.30	N/A
Site 1 128K	10/18/02 11:35:09	128	0.00888	76	76	0	37.80	29.70	N/A
Site 2 256K	10/18/02 11:35:09	76	0.01940	256	256	0	35.50	23.30	N/A
Site 3 512K	10/18/02 11:35:09	1518	0.03112	22	22	0	0.00	157.70	N/A
Site 4 768K	10/18/02 11:35:09	256	0.04968	226	226	0	48.40	41.90	N/A
Site 5 1M	10/18/02 11:35:09	909	0.06132	86	86	0	49.90	107.40	N/A
Site 6 128K	10/18/02 11:35:09	812	0.00787	11	11	0	53.30	94.40	N/A
Site 7 256K	10/18/02 11:35:09	1230	0.01561	11	11	0	38.50	140.80	N/A
Site 8 512K	10/18/02 11:35:09	430	0.03215	86	86	0	55.90	60.00	N/A
Site 9 768K	10/18/02 11:35:09	512	0.04788	116	116	0	59.70	68.90	N/A
Site 10 1M	10/18/02 11:35:09	1118	0.06107	66	66	0	48.10	120.70	N/A
Total	10/18/02 11:35:35	N/A	80.00000	1,275	1,275	0	62.00	22.10	N/A
A Group	10/18/02 11:35:35	N/A	N/A	1,275	1,275	0	62.00	22.10	N/A
Site 1 128K	10/18/02 11:35:35	128	0.01184	105	105	0	34.70	27.20	N/A
Site 2 256K	10/18/02 11:35:35	76	0.02587	334	334	0	36.00	22.10	N/A
Site 3 512K	10/18/02 11:35:35	1518	0.04150	33	33	0	0.00	153.60	N/A
Site 4 768K	10/18/02 11:35:35	256	0.06624	304	304	0	48.10	40.90	N/A
Site 5 1M	10/18/02 11:35:35	909	0.08176	114	114	0	49.90	102.60	N/A
Site 6 128K	10/18/02 11:35:35	812	0.01050	11	11	0	55.60	99.00	N/A
Site 7 256K	10/18/02 11:35:35	1230	0.02082	22	22	0	46.50	131.50	N/A
Site 8 512K	10/18/02 11:35:35	430	0.04286	114	114	0	55.00	60.70	N/A
Site 9 768K	10/18/02 11:35:35	512	0.06384	154	154	0	59.00	68.80	N/A
Site 10 1M	10/18/02 11:35:35	1118	0.08143	84	84	0	47.90	121.70	N/A
Total	10/18/02 11:35:59	N/A	100.00000	1,594	1,594	0	62.30	22.30	N/A

Figure 8-32. How Rates per Flow Affect Test Results (Jumbo Test Detail Example)

The test in this example also used various frame sizes for each flow, which can also be used with rates per flow.

One of the ways the Flow Rate Calculator displays the rate per flow is in frames per second.

Note: Throughput summary test results also allow you to view the rate as frames per second. To view the rate as frames per second, right-click on the test results and select **Frames/Sec** from the drop-down menu.

Figure 8-33 on page 275 shows an example of the throughput for each flow in frames per second.

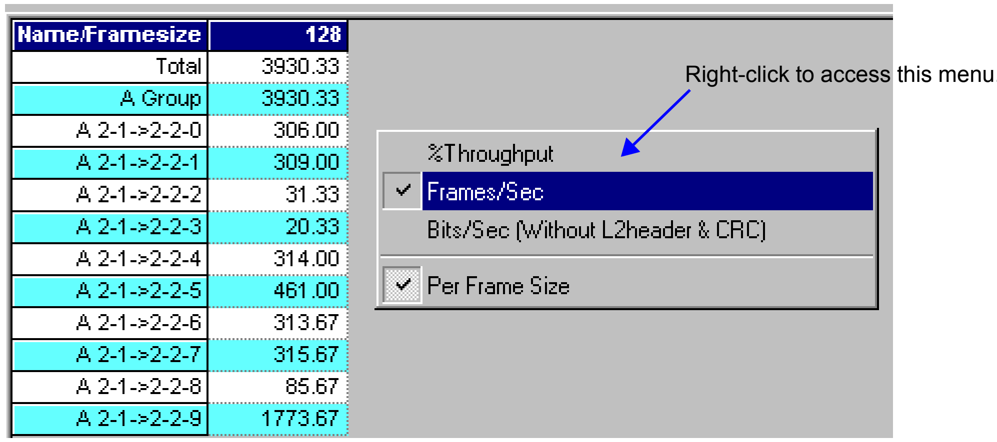

Figure 8-33. Throughput Test Summary Results Viewed in Frames per Second

Rates per Flow and per Iteration

When you vary the load (rate) *per flow and per iteration* (as opposed to per flow only or per port), the rate you assign to a flow determines the proportion of port bandwidth for the flow in relation to other flows on that port for a particular test iteration. The flow's proportion of port load can change for each test iteration.

For example, a flow using a rate sequence for rates per flow and iteration may look like this:

Iteration 1 - Flow1's rate is 30% of available port load
Iteration 2 - Flow1's rate is 20% of available port load
Iteration 3 - Flow1's rate is 20% of available port load

Different flows can have the same set of rates associated with them. Flows with rates set per flow and iteration are an uneven, user-defined distribution on the port. By contrast, when you do not assign rates per flow and iteration, each flow on the port sends roughly the same number of frames per second (based on the settings on the *Test Iterations* tab).

You can assign rates for flows on a port by specific percentages, or frame or bit rate units (such as fps or Kbps). Varying test load per flow and iteration is also known as custom loads per flow. Use the *Custom Loads Per Flow Table* dialog box on the *Test Setup>Test Iterations* tab to create a sequence (set) of custom rates that a flow will use for each iteration. Then assign a set of custom loads to a flow on the *Groups* tab.

Notes:

- The rates per flow and iteration feature is only available for LAN-33xx, XLW-372x, and XFP-373x modules.
- If a test contains more iterations than there are rates in a sequence, the last rate is reused for the remaining test iterations.

You can use rates per flow and iteration with custom test loads, but not with true random test loads. You can use frame size automation and custom frame sizes per flow with rates per flow when the rate unit is a percentage, but not with true random frame sizes.

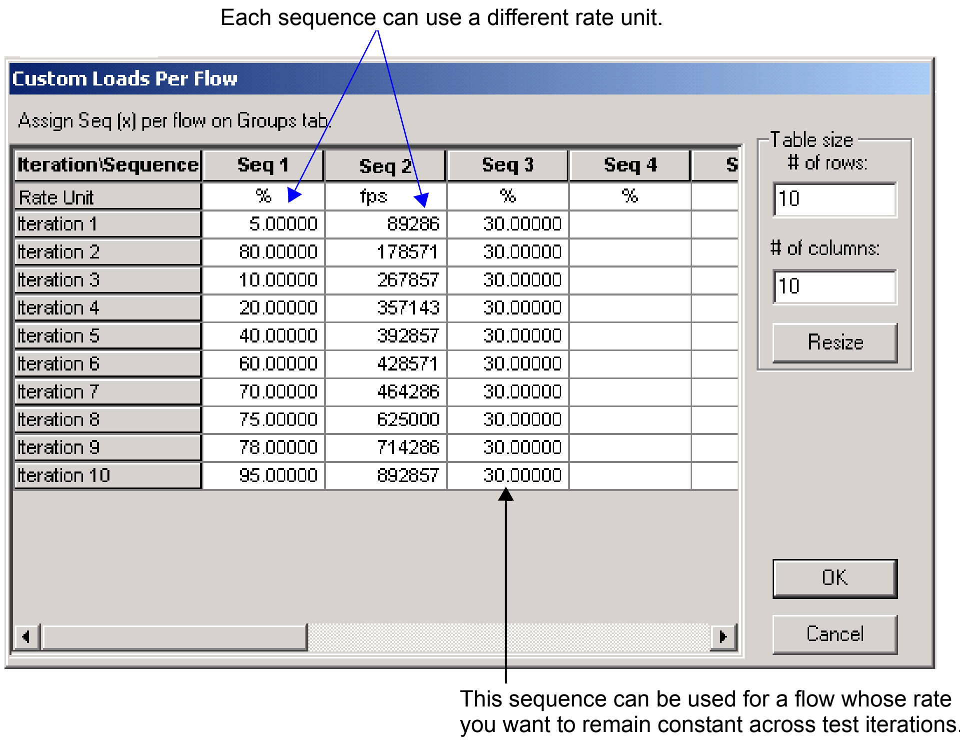

Figure 8-34. Custom Loads per Flow Table

For more detailed information on varying test load per flow and iteration, refer to the online Help.

Throttling Flows

For each test iteration, a flow's rate can increase, decrease, or remain constant while other flow rates change. You can throttle flows by increasing the rate of some flows with each new test iteration, while decreasing the rate of others and holding the rate of others constant. This scenario is useful for simulating increasing and decreasing network bandwidth conditions. Time sensitive data like voice and video are good candidates for this type of testing.

Example:

In this example, Seq3 has the same rate for each iteration. This will keep the rate of any flow using this sequence constant across iterations.

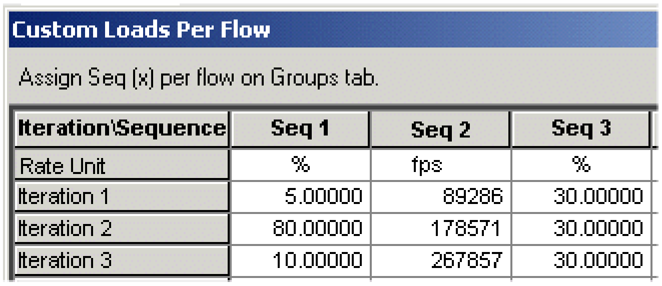
Custom Loads Per Flow

Assign Seq (x) per flow on Groups tab.

Iteration\Sequence	Seq 1	Seq 2	Seq 3
Rate Unit	%	fps	%
Iteration 1	5.00000	89286	30.00000
Iteration 2	80.00000	178571	30.00000
Iteration 3	10.00000	267857	30.00000

How Rates per Flow and Iteration Works

When you assign a rate per flow and iteration, each flow’s percentage of the port load can vary with each test iteration, regardless of the unit used to specify the percentage. The flow's proportion of the port bandwidth can change each time the test load increases with a new test iteration.

The rate percentage of each flow varies only after an iteration. For example, assume that Flow 1 uses a custom load sequence consisting of 20%, 30% and so on. The flow rate percentage for Iteration 1 will be 20% for the entire iteration. For Iteration 2, the flow's rate percentage will change to 30%, and so on.

Cyclic Flows:
When you assign a rate to a cyclic flow, that rate is evenly distributed among all of the cyclic variations. You cannot assign a rate for each variation of a cyclic flow, but the rate you assign applies to the base flow.

Rates per Flow and Iteration When Specified as a Percentage

On the *Custom Loads Per Flow* dialog box on the *Test Setup>Test Iterations* tab, when you specify a rate for a flow and iteration in percentage, the rate for the flow in that iteration will be a percentage of the port's available bandwidth, based on the test iteration rate.

For example, assume the following:

Flow 1 uses custom loads per flow Seq1, which contains Iteration 1 - 20%.
Test Iteration 1 is at 40%.
There are no other flows on the Port A.

Flow 1 is transmitted at 8% of Port A's bandwidth.

When There are More Test Iterations than Rates in a Sequence

SmartFlow uses the last rate specified in the sequence for any additional test iterations. If it is in a bit or frame rate unit, SmartFlow converts the last specified unit to a percentage. As test load changes, that percentage rate is no longer equivalent to the specified number of Kbps for the last iteration in the sequence. It will vary according to the test load for that iteration and the port's available bandwidth in that iteration.

For example, on a 100 Mbps port, 600 Kbps at a 40% test load is 1.5% of the port speed. If the test load increases to 70%, 1.5% for the flow rate now becomes 1050 Kbps.

Rates per Flow and Iteration When Specified in Rate Units

The calculation for Mbps, Kbps and bps is based on:

Preamble + Layer 2 Header + Payload +CRC + IFG

The IFG is fixed at 96 bits.

When you specify a rate for a flow and iteration in the *Custom Loads Per Flow* dialog box (on the *Test Setup>Test Iterations* tab) as a rate unit, the flow is transmitted at that exact rate specified as long as there is available bandwidth on the port in that iteration. Normally bandwidth validation for ports occurs on the *Groups* tab.

Note: If there are more test iterations than rates specified in the *Custom Loads Per Flow* dialog box and if the rates are in units (e.g., fps and Mbps), the transmission rate for the subsequent "extra" iterations is not the same as the rate in units for the last iteration specified in the table. This is because SmartFlow converts the last rate specified into a percentage for the extra (unspecified rate) test iterations. Using a percentage means the flow is transmitted at a proportion of the current test load for the remaining test iterations.

Examples of Rates per Flow and Iteration

The rates per flow and iteration feature is best understood by examples.

Example 1: Fewer Rates Specified Than Actual Test Iterations, *page 279*.

Example 2: Rates per Flow and Iteration Bit Rate Units and Available Bandwidth, *page 279*.

Example 3: Frame Size Automation Used with Rates per Flow and Iteration, *page 280*.

Example 4: Rates per Flow Used with Rates per Flow and Iteration, *page 281*.

Example 5: IMIX used with Rates per Flow and Iteration, *page 283*.

Example 1: Fewer Rates Specified Than Actual Test Iterations

Assume that the following four rates were set up in the *Custom Loads Per Flow* dialog box but the test has seven iterations:

Custom Test Load Table Iteration	Rate	Test Iteration
1	23.400000	**1**
2	40.500000	**2**
3	80.300000	**3**
4	40.000000	**4**
	40.000000	**5**
	40.000000	**6**
	40.000000	**7**

Example 2: Rates per Flow and Iteration Bit Rate Units and Available Bandwidth

Assume that four rate iterations were specified in the *Custom Test Loads Per Port (% Utilization)* dialog box, with 600 Kbps in the last iteration (row).

Port speed is 100 Mbps and there is only one flow on the port.

Assume there are eight test iterations, and this is four more than what was specified in the *Custom Test Loads Per Port (% Utilization)* dialog box. SmartFlow uses the last rate specified in the table for any additional test iterations. If it is in a bit or frame rate unit, SmartFlow converts the last specified unit to a percentage. In this case, 600 Kbps is equivalent to 1.5% of the test load for the last iteration specified. This means that 1.5% is used as the intended flow rate for the remaining “unspecified” test iterations.

However, as test load increases, 1.5% is no longer equivalent to 600 Kbps. It varies according to the test load for that iteration and the port's available bandwidth in that iteration. If the port speed is 100 Mbps and the test load increases to 70%, 600 Kbps is now only 1.5% for the flow rate or 1050 Kbps.

SmartFlow calculates the flow rate percentage for the last specified iteration in the sequence by:

Iteration 4: 600 Kbps/40 Mbps = 1.5%

SmartFlow calculates the flow rate percentage for subsequent iterations by:

Iteration 5: 1.5% of 50 Mbps = 750 Kbps

Gray rows represent test iterations that were not specified in the *Custom Test Loads Per Port (% Utilization)* dialog box sequence.

Test Iteration	Seq1 Rate in Kbps	Test Iteration Rate	Test Iteration Rate in Mbps
1	500	10%	10 Mbps (10% of 100 Mbps)
2	300	20	20 Mbps (20% of 100 Mbps)
3	200	30	30 Mbps (30% of 100 Mbps)
4	600	40	40 Mbps (40% of 100 Mbps)
5	750	50	50 Mbps (50% of 100 Mbps)
6	900	60	60 Mbps (50% of 100 Mbps)
7	1050	70	70 Mbps or 1050 Kbps (50% of 100 Mbps)
8	1200	80	80 Mbps or 1200 Kbps (50% of 100 Mbps)

Example 3: Frame Size Automation Used with Rates per Flow and Iteration

Assume that you have set up the following frame sizes in the *Custom Frame Sizes (all ports)* dialog box: 64, 128, 512.

Assume that you set up a load sequence (Seq1) for rates per flow and per iteration consisting of: 25%, 35%, 45%.

The test is set up to step up the load by 10% with a minimum load of 10% and maximum load of 30%, which would yield three test iterations.

In frame size automation, the test essentially starts over for each new frame. Therefore, even though the *Custom Loads Per Flow* dialog box contains only three iterations, the Seq1 Iteration 1 rate (25%) is used again for test iteration #4, Seq 1 Iteration 2 rate (35%) is used again for test iteration #5, and so on until the entire test completes.

Test Iteration	Frame Size	Test Iteration Rate	Seq1 Rate	Intended % Rate for Flow 1
1	64	10%	25%	2.500000%
2	64	20%	35%	7.000000%
3	64	30%	45%	13.500000%
4	128	10%	25%	2.500000%
5	128	20%	35%	7.000000%
6	128	30%	45%	13.500000%
7	512	10%	25%	2.500000%
8	512	20%	35%	7.000000%
9	512	30%	45%	13.500000%

Example 4: Rates per Flow Used with Rates per Flow and Iteration

Assume the following:

There are four flows on Port 1: Flow1, Flow2, Flow3, Flow4

Flow1 and Flow2 use rate per flow and per iteration flows:

Flow1 = Seq1
Flow2 = Seq2

Seq1: 20%, 40%, 20%, 35%, 15%, 10%
Seq2: 10%, 40%, 30%, 30%, 30%, 30%

Flow3 = 15%
Flow4 = even distribution

Flow1 and Flow2's rates are based on the appropriate rate in Seq1 and Seq2, respectively, and so will vary with each test iteration. Flow3's rate is per flow only (not also per iteration), so its rate of 15% remains constant across iterations. Flow4's rate is “even distribution”, which is also on a per flow basis. However, it varies according to the remaining bandwidth for the iteration.

As the test progresses, SmartFlow transmits flows at the rates shown in this table:

Flow	Displayed Flow Rate (%/Rate unit/Custom column on the Groups tab)	% Rate for Flow
Test Iteration 1		
Flow 1	Seq1	20%
Flow 2	Seq2	10%
Flow 3	15%	15%
Flow 4	Even Distribution	55% (100- [20+10+15])
Test Iteration 2		
Flow 1		
Flow 2		
Flow 3		
Flow 4		
Test Iteration 4		
Flow 1	Seq1	35%
Flow 2	Seq2	30%
Flow 3	15%	15%
Flow 4	Even Distribution	20% (100- [30+35+15])

Example 5: IMIX Used with Rates per Flow and Iteration

Assume you used the following IMIX set of weighted frame sizes for a set of flows:

Frame Size	Packet %
56	50%
1518	25%
R	25%

Assume that you also used for rates per flow and iteration plus the Custom Test Loads table contains these sequences:
Seq1: 10%, 15%, 50%, 75%, 100%
Seq2: 10%, 10%, 10%, 10%, 10%

The test setup includes a stepped test load: minimum 10%, step 10%, maximum 50%.

Flow	Displayed Flow Rate (%/Rate unit/Custom column on the Groups tab)	Intended % Rate for Flow
Test Iteration 1 (10%)		
Flow 1	Seq1	1% (10% of 10% test load)
Flow 2	IMIX (frame size 56 - 50%)	3% for total of IMIX flows
Flow 3	IMIX (frame size 1518 - 25%)	
Flow 4	IMIX (random frame size - 25%)	
Flow 5	20% (rate per flow)	2% (20% of 10% test load)
Flow 6	Even distribution	3%
Flow 7	Seq2	1% (10% of 10% test load)

Flow	Displayed Flow Rate (%/Rate unit/Custom column on the Groups tab)	Intended % Rate for Flow
Test Iteration 2 (20%)		
Flow 1	Seq1	3% (15% of 20% test load)
Flow 2	IMIX (frame size 56 - 50%)	6.5% for Total of IMIX flows
Flow 3	IMIX (frame size 1518 - 25%)	
Flow 4	IMIX (random frame size - 25%)	
Flow 5	20% (rate per flow)	4% (20% of 20% test load)
Flow 6	Even Distribution	6.5% (20-7=13/2)
Flow 7	Seq2	2% (10% of 20% test load)

Specifying a Custom Loads per Flow and Iteration Sequence

The *Custom Loads Per Flow* dialog box on the *Test Setup>Test Iterations* tab allows you to specify a predefined series of rates known as a sequence that a flow should use at each iteration of a test. This dialog box is used when you vary rates per flow and per iteration (as opposed to per flow only), where you assign a sequence of rates that a flow will use at each test iteration. The rates per flow and per iteration feature is only available for LAN-33xx, XLW-372x, and XFP-373x modules.

Figure 8-35 illustrates how flows using custom loads per flow and iteration appear on the *Groups* tab.

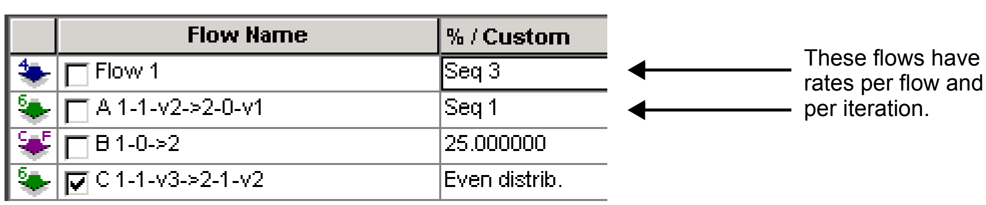

Figure 8-35. How Flows Using Custom Loads per Flow and Iteration Appear on Groups Tab

You assign a sequence to a flow at the *Groups* tab in the *%/Rate unit /Custom* column by selecting one from the drop-down list, as shown in *Figure 8-36*.

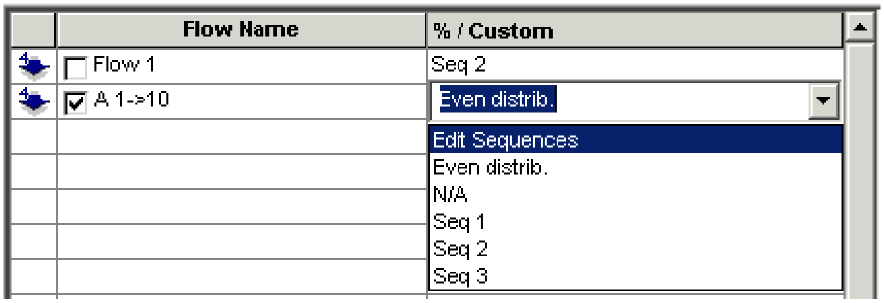

Figure 8-36. %/Rate Unit/Custom Drop-down Box

When setting up rate sequences in the *Custom Loads Per Flow* dialog box on the *Test Setup>Test Iterations* tab, if you change the rate unit displayed from a percentage to a frame rate or bit rate (Mbps, kbps, bps) or vice versa, the *Display Rate Conversion* dialog box is displayed.

For more detailed information about how to specify loads per flow and iteration, refer to the online Help.

Important: If you change the number of rows or columns in the *Custom Loads Per Flow* dialog box, make sure you click the **Resize** button for the changes to take effect.

Blank sequences can exist. However, all sequences must be either all blank or all filled.

Notes:

- If a rate sequence contains fewer rates than there are test iterations, the last rate is reused for the remaining test iterations.
- If a rate sequence contains more rates than test iterations (for example, six table rows and only four test iterations), the extra table values are not used in the test.
- Custom loads per flow and iteration only applies to a Throughput test when percentages are used for rates since it uses a set number of test iterations, so the overall load cannot be determined.
- If the rate for a flow is less than 1 fps, test results may be invalid.

Creating a Constant Rate for a Flow

If you want to create a constant flow rate where one flow remains at a constant bit or frame rate while other flow rates increase or decrease at each iteration, in a new load sequence enter the same rate for all cells in the appropriate sequence. If you simply want the same rate (proportion of bandwidth) for a flow across iterations, use a percentage instead of a bit or frame rate.

How Bit or Frame Rate Units are Calculated

Bit or frame rate units are calculated into a percentage of 100% of port load. For example, 10,000 Kbps when the test load is 10% is 1% for a gigabit port. As long as there is bandwidth available on the port for the flow at that rate, the flow transmits at that rate, regardless of the test iteration rate. If there is not enough bandwidth, you will get a validation error on the *Groups* tab.

When you change the rate unit between a percentage and a frame rate or bit rate (Mbps, kbps, bps), the *Display Rate Conversion* dialog box is displayed. This dialog box allows you to select the port rate and/or frame size, so that you can see the actual rate in the unit that you selected at which the flow will be transmitted.

How Rates per Flow and Iteration Affects Test Results

When you set up a test with rates assigned per flow and iteration, test results display each flow with a different rate within the same test iteration and a different rate across iterations.

Figure 8-37 shows an example of results with flows using rates per flow and iteration.

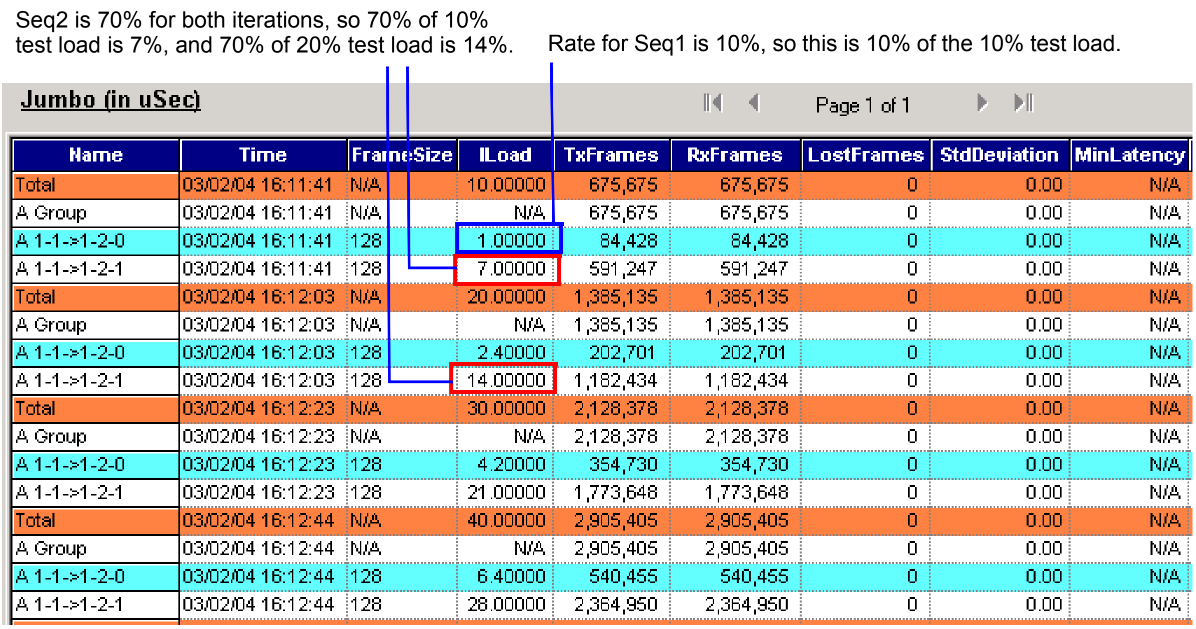

Name	Time	FrameSize	ILoad	TxFrames	RxFrames	LostFrames	StdDeviation	MinLatency
Total	03/02/04 16:11:41	N/A	10.00000	675,675	675,675	0	0.00	N/A
A Group	03/02/04 16:11:41	N/A	N/A	675,675	675,675	0	0.00	N/A
A 1-1->1-2-0	03/02/04 16:11:41	128	1.00000	84,428	84,428	0	0.00	N/A
A 1-1->1-2-1	03/02/04 16:11:41	128	7.00000	591,247	591,247	0	0.00	N/A
Total	03/02/04 16:12:03	N/A	20.00000	1,385,135	1,385,135	0	0.00	N/A
A Group	03/02/04 16:12:03	N/A	N/A	1,385,135	1,385,135	0	0.00	N/A
A 1-1->1-2-0	03/02/04 16:12:03	128	2.40000	202,701	202,701	0	0.00	N/A
A 1-1->1-2-1	03/02/04 16:12:03	128	14.00000	1,182,434	1,182,434	0	0.00	N/A
Total	03/02/04 16:12:23	N/A	30.00000	2,128,378	2,128,378	0	0.00	N/A
A Group	03/02/04 16:12:23	N/A	N/A	2,128,378	2,128,378	0	0.00	N/A
A 1-1->1-2-0	03/02/04 16:12:23	128	4.20000	354,730	354,730	0	0.00	N/A
A 1-1->1-2-1	03/02/04 16:12:23	128	21.00000	1,773,648	1,773,648	0	0.00	N/A
Total	03/02/04 16:12:44	N/A	40.00000	2,905,405	2,905,405	0	0.00	N/A
A Group	03/02/04 16:12:44	N/A	N/A	2,905,405	2,905,405	0	0.00	N/A
A 1-1->1-2-0	03/02/04 16:12:44	128	6.40000	540,455	540,455	0	0.00	N/A
A 1-1->1-2-1	03/02/04 16:12:44	128	28.00000	2,364,950	2,364,950	0	0.00	N/A

Figure 8-37. How Flows with Rates per Flow and Iteration Appear in Test Results

Minimum/Maximum Rates for Gigabit and POS Ports

When determining the minimum and maximum rates for tests involving Gigabit and POS ports, keep in mind the following limits for these cards in relation to the number of flows per port.

For more detailed information about maximum and minimum rates for Gigabit and POS ports, see the *SmartBits System Reference Manual* that accompanied your SmartBits chassis.

Minimum/Maximum Frames per Second

Gigabit/POS card minimum rate: 31 frames per second per flow

The formula to calculate the minimum rate allowed per Gigabit/POS port is as follows:

[(31 fps X no. flows per port) / maximum fps] X 100

The formula to calculate the maximum rate allowed per Gigabit/POS port is as follows:

[Maximum fps (that card can send per second depending on packet size) / no. flows per port] X 100

Note: The LAN-33xx Ethernet module allows a minimum rate of 1 frame per second per flow.

Setting up Global Application Preferences

You can specify global application preferences, such as whether to display flows as well as groups in test results. Global preferences affect all tests and their results. These options include:

- Whether to stop tests when errors occur
- Whether to show flows (as well as groups) in online results
- Whether to delay before transmitting or before reading the counter.

To specify other test parameters, such as duration and load parameters, use the *Test Setup* tab.

Specify application preferences on the *Options* tab of the main window. The *Options* tab contains two sub-tabs:

- *General* tab (*Figure 8-38*)
- *Results* tab (*Figure 8-39 on page 291*)

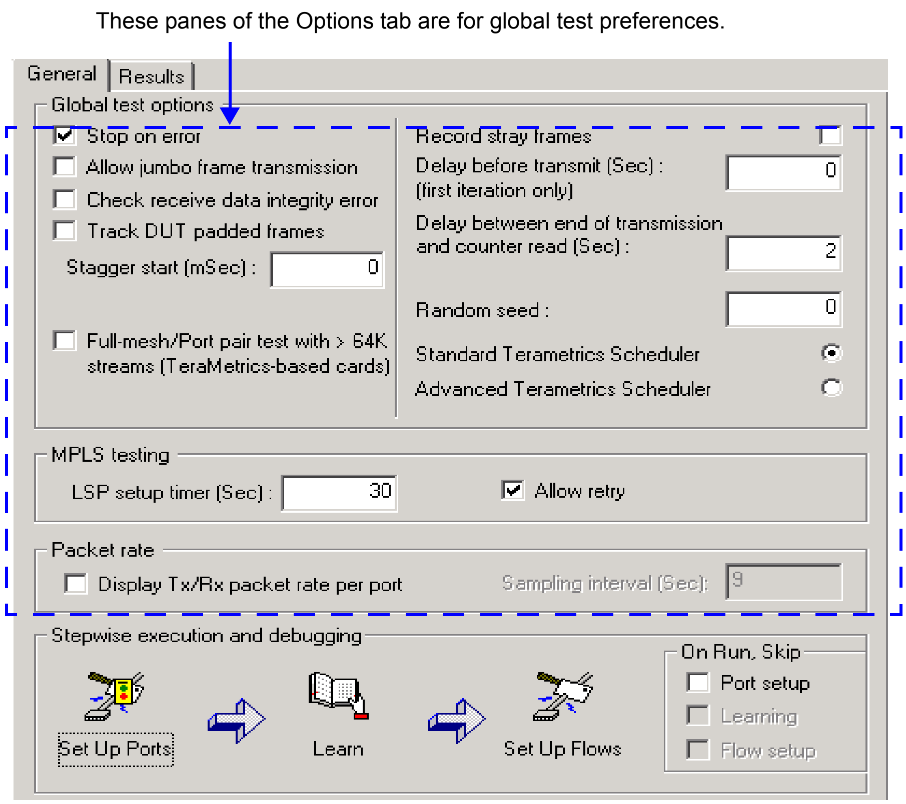

Figure 8-38. Global Application Preferences on the Options General Tab

None of the test preferences on this tab are required. For information about the *Stepwise execution and debugging* pane on the lower part of this tab, see *"Debugging a Test" on page 294*. For information about how delays affect mixed traffic, see *"How Delays Affect Mixed Traffic" on page 484*.

The *Options>General* tab allows you to:

- Set global application preferences for SmartFlow.
  Global preferences affect all tests and their results. For example, one option is whether to stop tests when errors occur.
- Choose the Standard TeraMetrics Scheduler or Advanced TeraMetrics Scheduler. (Default is Standard TeraMetrics Scheduler.)
- Set various other options including: MPLS testing options and packet rate options.
- Execute (manually) the phases of a test in steps (using the *Stepwise execution and debugging* pane).

To specify other test parameters, such as duration and load parameters, use the *Test Setup* tab.

Buttons and fields in the lower half of the tab relate to the step-by-step test execution.

Note: None of the fields on this tab are required.

Important: There are two versions of the TeraMetrics scheduler: the Standard (original) TeraMetrics Scheduler and the Advanced TeraMetrics Scheduler. The Standard TeraMetrics Scheduler may adjust requested Tx rates on a per-flow basis. This may be noticeable at higher link speeds. The Advanced TeraMetrics Scheduler uses a new rate-based algorithm that provides a better, more even distribution of streams. This new algorithm alleviates burstiness in the transmission of unequal stream.

Note: The Advanced TeraMetrics Scheduler is a test-wide option and is supported on the following cards: TeraMetrics-based XD, TeraMetrics-based POS, and the 10Gigabit modules. Other TeraMetrics-based modules support only the Standard TeraMetrics Scheduler.

For a description of each field on the *Options>General* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Global Test Options

The *Global test options* pane on the *Options>General* tab contains fields that affect tests globally.

For a description of each field on the *Options>General* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

MPLS Testing Options

The *MPLS testing* pane on the *Options>General* tab contains two fields used to customize MPLS testing.

For a description of each field on the *Options* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Packet Rate Options

Options in the *Packet rate* pane on the *Options>General* tab allow you to specify whether to display virtual real-time transmit and receive rates (in packets per second) for each port in the test results and for each iteration, in addition to the overall final rate of the port. Packet rate results appear on a separate *Packet Rate* test results tab, which can display up to 20 ports in a test. For information about the *Packet Rate* test results tab, see *"Packet Rate Tab" on page 83*.

When you set a sampling interval at which to measure packet rate, keep in mind that the shorter the time sampling, the more frequently SmartFlow counts packets and thereby uses more computer resources and thus affects performance. Resources are affected by:

- Time sampling interval
- Number of flows in the test
- PC memory.

For a description of each field on the *Options* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Setting Up Global Results Preferences

Specify application preferences on the *Options>Results* tab of the main window, as shown in *Figure 8-39 on page 291*.

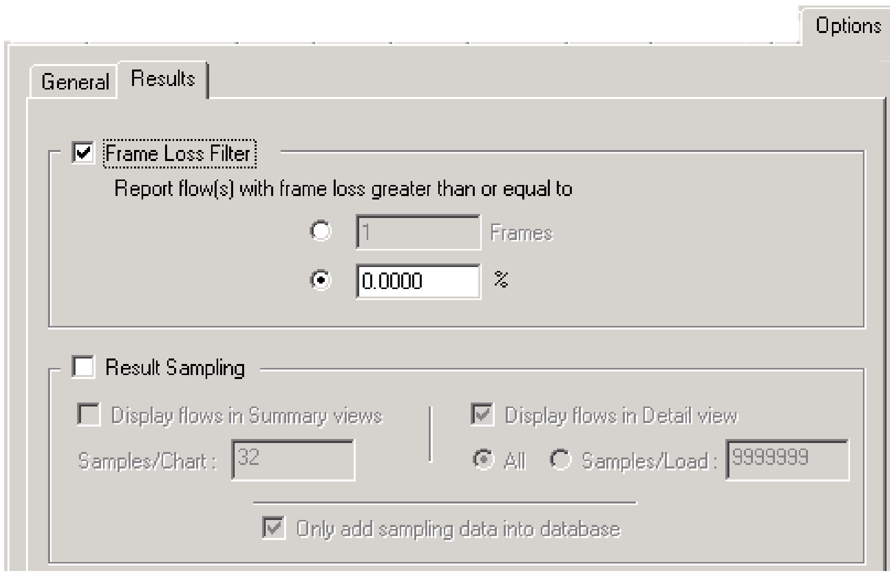

Figure 8-39. Global Results Preferences on the Options Results Tab

Use the *Options>Results* tab to:

- Set frame loss reporting options.
- Set results display and sampling.

Selecting the *Result Sampling* checkbox disables the *Frame Loss Filter* checkbox and associated options, and vice versa.

Results Frame Loss Filter Options

These options allow you to display only flows with frame loss in results. This is useful if you have a large configuration or run tests of long duration that produce a lot of results. These options allow you to focus on only the troublesome areas of your test.

Note: The *Frame Loss Filter* checkbox and associated options only apply to the Throughput, Jumbo, and Frame Loss tests. If these options are set but the test is not one of the allowed tests, SmartFlow uses the default settings for result sampling, which are the selected *Display flows in Detail view/All* options.

The total and group lines are always displayed in results, regardless of the frame loss filter option selected.

Results Display and Sampling Options

The *Result Sampling* pane on the *Options>Results* tab contains options that control whether flows (in addition to just groups and totals) should appear in results and how many samplings should be taken of test results. These options affect the detail level of results, and can also affect performance plus whether a test reaches the 1-Gigabyte size limit of the database that stores results. For more information about how to use these fields to adjust the size of results, see *"Test Setup Guidelines for Large Tests" on page 212*.

As an example of how these fields interrelate, assume you set up 189 flows and activate the results sampling options as shown in *Figure 8-40*.

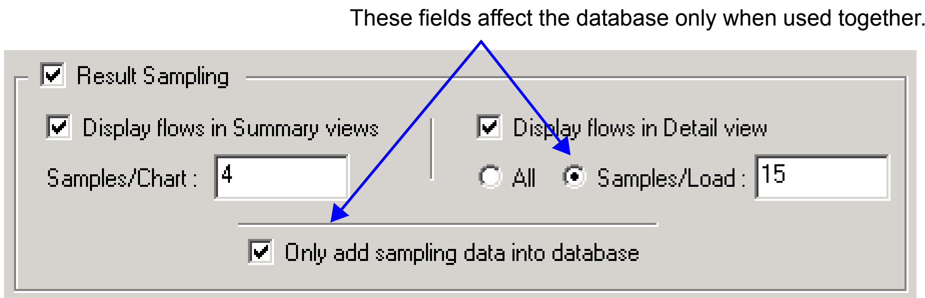

Figure 8-40. Example Results Sampling Setup

Based on these sampling options, SmartFlow only takes four samples from the 15 samples that were written to the database, not four samples from the 189 flows.

For a description of each field on the *Options>General* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Customizing Audible Test Alarms

SmartFlow provides audible alarms that mark stages of the test that is running. When SmartFlow is installed, it copies default sound files to the miscellaneous directory located wherever the SmartFlow application was installed. *Table 8-8* shows the SmartFlow events that can be associated with a sound file and that are listed in the *Sound Properties* dialog on your PC.

Table 8-8. SmartFlow Events

Event Name	Type	When Alarm Sounds
Started	Notification	This alarm sounds when the setup phase of the test completes. The setup phase is when SmartFlow verifies flows according to the port and network setup prior to transmission.
Stopped	Notification	This alarm sounds when the test completes.

You can turn off the sound or select another sound file for these SmartFlow events. Use the *Sound Properties* dialog box, as shown in *Figure 8-41*. For the procedure to change or turn off an alarm sound, refer to the online Help.

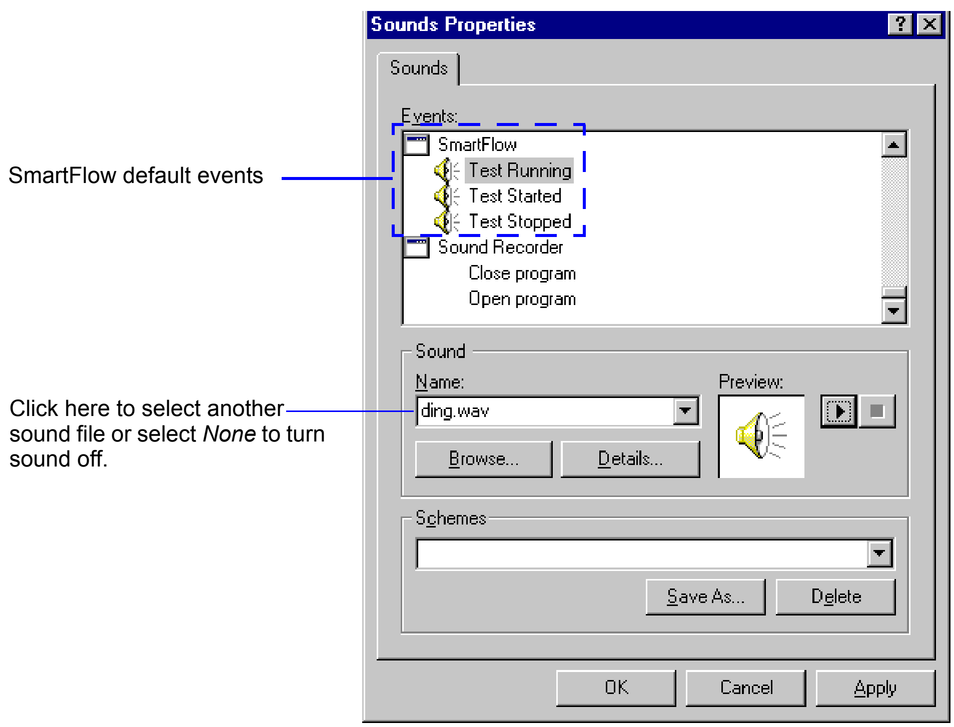

Figure 8-41. Sound Properties Dialog Box

Debugging a Test

SmartFlow contains some debugging features to help determine what may be preventing a test from running. You can also use these features to manually execute the phases of a test in steps, rather than having it run automatically when you launch a test. By isolating these phases, you have more control over the test and can prevent a test from failing to run. These options are available on the *Options* tab (*Figure 8-42*) in the main window.

The *Stepwise execution and debugging* pane on the *Options>General* tab allows you to run only the learning and/or setup phases of a test prior to running the test. Once you run the learning phase and/or the setup phases, you can launch a test from the shortcut bar.

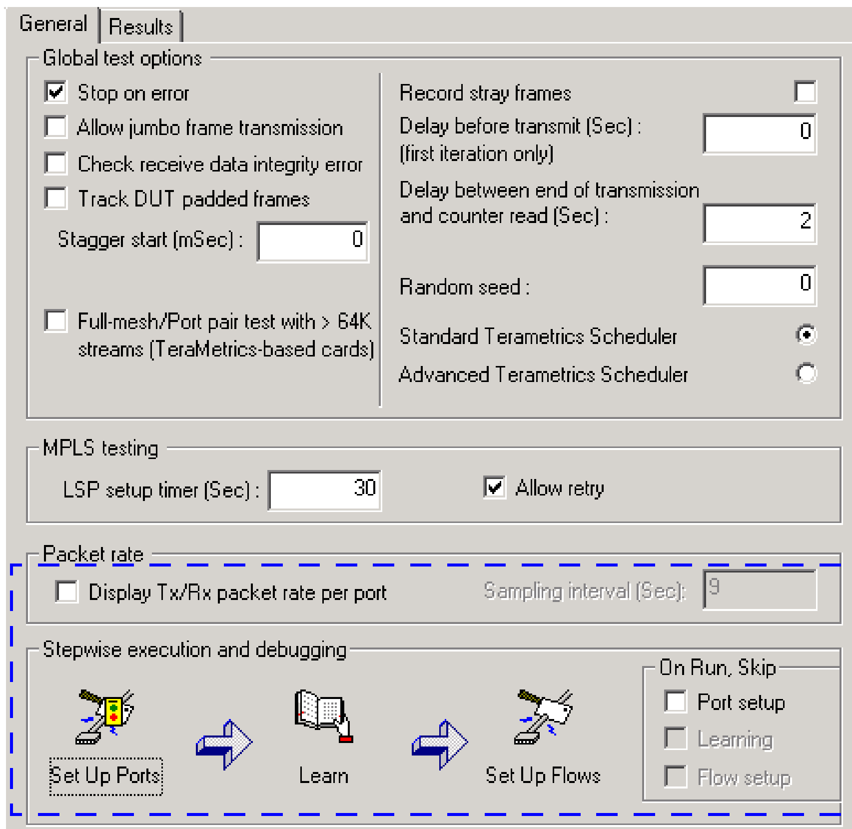

Figure 8-42. Using the Options tab to Debug a Test

For information about the options in the upper portion of the *Options* tab, see *"Setting up Global Application Preferences" on page 288*.

Skipping Test Phases When You Run the Test

When you run a test in individual steps, you may not want to run them again when you launch the test. To skip a test phase when you run the test, you must skip the previous step also. For example, if you want to skip the learning phase, you must also skip the port setup phase. For more information about skipping a test phase, see the individual test phase description in the following paragraphs.

Running the Port Setup Phase

Once you set up the test ports and connect to the SmartBits chassis, you can perform only the setup phase prior to running a test. During this phase, SmartFlow establishes communication between the DUT and the SmartBits chassis, downloads the following information to SmartBits cards, and then applies these settings before transmission occurs:

- Port setup (such as speed)
- Test setup
- Flow setup.

Note: You do not have to run the learning phase prior to running the download setup step.

Use the *Set Up Ports* (download setup) stepwise execution to:

- Check port setup before a test starts running and spot any problem that may cause the test to fail as a result.
  For example, you can see if the DUT failed to autonegotiate to the speed of the SmartBits ports.
- Provide the DUT time to set the autonegotiation to the speed of the SmartBits ports. For example, in order for the DUT to set its speed to 100 Mbps for SmartBits ports with this speed, in a normal test the DUT may not have time to do so before the ports start transmitting test packets at this speed.

To perform the download setup phase:

1. Click the **Options>General** tab.
2. Click the **Setup Ports** button.

Once the download step is complete, the status appears at the bottom of the window:

Skipping the Port Setup Phase During Tests

If you already ran the learning phase and the port setup phase in stepwise executions, you may want to skip the port setup when you run a test.

To skip the port setup phase during a test:

1. Click the **Options>General** tab.
2. Select the **Port setup** checkbox in the **On Run, Skip** pane.

3 Go to the shortcut bar and click a test launch button to start a test, or go to the section *"Running the Learning Phase" on page 296* to perform learning phase next.

Running the Learning Phase

The learning phase of a test takes place prior to the transmission of any test packets. If this phase fails, a message box appears showing the port that failed and prompts you for further action. However, you can also perform learning as a separate step, independent of the actual test.

Use the stepwise execution of the learning phase to:

- Prevent failures due to learning.

  For example, if you plan to run a test overnight and you do not want it to fail due to a learning error, you can first use the stepwise execution of learning. A stepwise execution of the learning phase lets you correct any learning failures prior to running actual tests.
- Save time by running learning only once for multiple tests.

Note: The independent execution of the learning phase does not eliminate the learning phase from tests that you launch from the shortcut bar. If you perform the learning step independently and then launch the test, the learning phase still occurs prior to sending out test packets.

Once you connect to the SmartBits chassis, configure all of the test ports and specify the learning type (Layer 2 MAC address learning or Layer 3 ARP requests) on the *Cards* tab.

To perform the learning phase:

1 Click the **Options>General** tab.

2 Click the **Learn** button.

Once the learning step is complete, the status appears at the bottom of the window:

Go to the shortcut bar and click a test launch button to start a test, or go to *"Running the Flow Setup Phase" on page 297* to run the flow setup phase next.

Skipping the Learning Phase During Tests

If you already ran the learning phase in a stepwise execution, you may want to skip it when you run a test.

To skip the learning phase during a test:

1 Click the **Options>General** tab.

2 Select the **Port setup** and **Learning** checkboxes in the **On Run, Skip** pane.

If you chose to skip learning, when the test runs you will still see a status message stating:
`x ports have resolved ARPs in testing phase.`
The purpose of this message is to learn the destination MAC so that the frame can reach the destination.

Running the Flow Setup Phase

The flow setup phase consists of downloading flow and group information specified in the Group Wizard and/or on the *SmartFlows* tabs. The flow setup phase takes place after the learning phase (if specified) but before the test starts running. You can perform this phase as a separate step, independent of the actual test.

Go to the shortcut bar and click a test launch button to start a test.

Skipping the Flow Setup Phase During Tests

If you already ran the flow setup phase in a stepwise execution, you may want to skip it when you run a test.

To skip the flow setup phase during a test:

1 Click the **Options>General** tab.

2 Select the **Port setup**, **Learning**, and **Flow setup** checkboxes in the **On Run, Skip** pane.

Note: You can only skip the flow setup phase of the test if you also skip port setup and learning.

IPv6 Test Setup

When implementing IPv6, typically you first measure the data plane performance (throughput, latency, frame loss, and QoS tests) of the DUT. Based on test results, adjust the DUT control plane, and then test the data plane again.

The main aspects of IPv6 that you may want to test are:

- *Raw IPv6 performance*
  Use this type of testing if only a single stack is used. See *"Raw IPv6 Performance" on page 299*.
- *Dual Stacks* (IPv4 and IPv6 traffic)
  Use this type of testing to determine if the DUT can handle both types of traffic and whether the IPv4 traffic suffers when IPv6 traffic is added. See *"Dual Stacks" on page 299*.
- *Tunneling* (IPv6 to IPv4)
  This method allows you to send IPv6 packets across an IPv4 infrastructure to "islands" of isolated IPv6 networks without changing any of the IPv4 components. See *"Tunneling" on page 300*.
- *Network Address Translation* (NAT) (IPv6 to IPv4, IPv4 to IPv6)
  This method, also known as Network Address Protocol Translation (NAPT), allows IPv6 hosts to communicate with IPv4 hosts and vice versa.

  A NAT device, sometimes known as an Application Level Gateway (ALG), is used that converts the IPv6 addresses to IPv4 addresses (and vice versa) by mapping one to the other. The address assignments can be done statically or dynamically. See *"Network Address Translation (NAT-PT)" on page 302*.

Recommended Testing Order

When testing IPv6, you may want to run tests that focus on the different aspects of IPv6 in this order:

1. *Get an IPv4 benchmark.*
   Run tests with only IPv4 traffic to get a benchmark of how the DUT handles only IPv4 traffic. Use this to compare with IPv6 test results.
2. *Test raw IPv6 performance alone.*
   Run tests with only IPv6 traffic to test how the DUT handles IPv6 alone without IPv4. Compare the IPv6-only test results with the IPv4-only test results.
3. *Test the DUT dual IPv4/IPv6 stack.*
   Run tests with both IPv4 and IPv6 traffic to see how the DUT handles IPv6 traffic in the presence of IPv4 traffic.

   For example, DUTs in which IPv4 traffic has a significant impact on even light IPv6 loads may not be suitable in an IPv6/IPv4 production environment.

4. *Test the DUT ability to tunnel IPv6 traffic over IPv4 networks.*
   Run tests with both IPv4 and IPv6 traffic to see how well the DUT sets up IPv4 tunnels and forwards IPv6 traffic over them.
5. Test the DUT ability to translate IPv6 packet to IPv4 packets and vice versa.

You can also test QoS during any of these test stages. See *"QoS Testing" on page 304*.

IPv6 Test Setup Tips

Set up tests for IPv6 traffic in the same way as other types of testing, with a few exceptions.

If you are testing a router with an IPv6 stack but IPv6 has not been implemented to use the switch fabric, the IPv6 traffic will exhibit lower throughput.

For any Layer 2 testing with IPv6, you must clear the *Addr Resolution* field on the *Cards* tab in order for neighbor discovery to take place.

Important: When you set up tests for tunneling and address translation, both the source and destination IP addresses of the flows must be IPv6 addresses, even if you are configuring an IPv6-to-IPv4 test or an IPv4-translated-to-IPv6 test.

Raw IPv6 Performance

This type of testing uses only a single stack and can include host-to-host, router-to-router, router-to-host, and host-to-router. Testing IPv6 with QoS may be included in this category. You can use this type of testing as a benchmark before beginning to test the more complex aspects of IPv6. *Figure 8-43* illustrates an example of host-to-host setup to test IPv6.

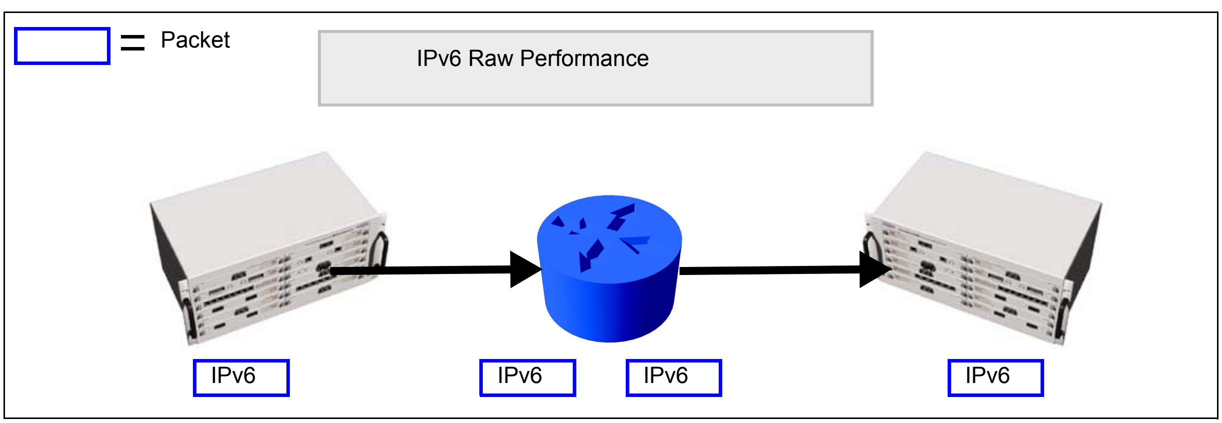

Figure 8-43. IPv6 Raw Performance Testing

Dual Stacks

In dual stack testing, SmartBits chassis ports can support both IPv6 and IPv4 flows to simulate a dual-stack mode and send packets to a dual stack DUT.

Figure 8-44 illustrates an example of host-to-host setup to test IPv6 and IPv4.

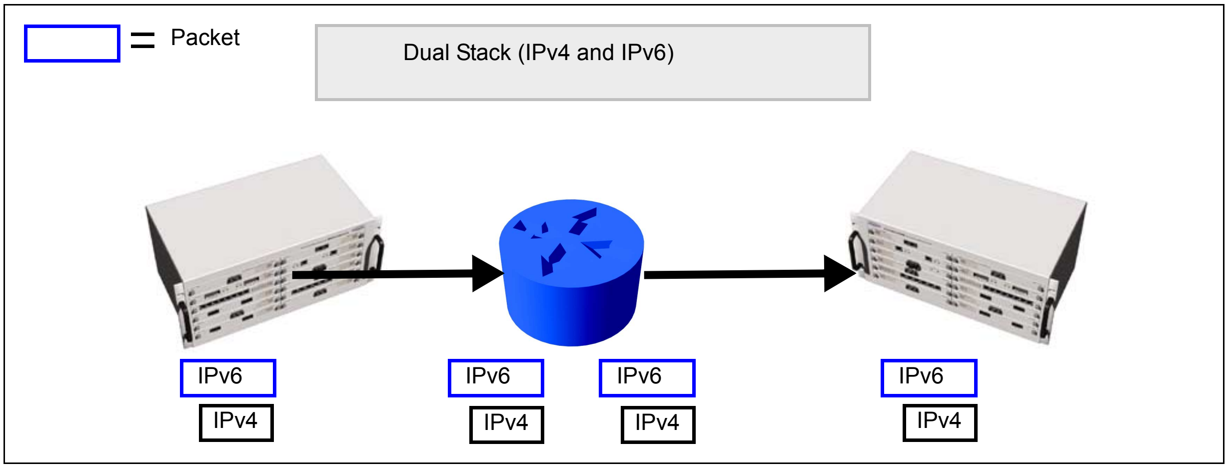

Figure 8-44. Dual Stack Performance Testing

Be sure to configure both IPv4 and IPv6 on each DUT interface that will be used in the test.

Tunneling

Tunneling allows you to send IPv6 packets across an IPv4 infrastructure to “islands” of isolated IPv6 networks without changing any of the IPv4 components. A border router encapsulates IPv6 packets in IPv4 packets and sends them across an IPv4 infrastructure. The router on the other side of the border then de-encapsulates the IPv4 packets back to IPv6 and routes them as IPv6 traffic to the destination.

With tunneling, SmartFlow transmitting ports simulate IPv6 transmitters (end nodes or border routers on the network), and the receiving nodes can be either IPv4 or IPv6. However, SmartFlow transmitting ports do *not* simulate IPv4 transmitters whose packets have IPv6 encapsulated within an IPv4 header.

Note: Do not confuse tunneling mechanisms (such as 6to4) with IPv6 to IPv4 address translation mechanisms.

Although other tunneling methods exist, SmartFlow has only been tested with the 6to4 method.

Figure 8-45 illustrates an example of host-to-host setup to test IPv6-to-IPv4 tunneling.

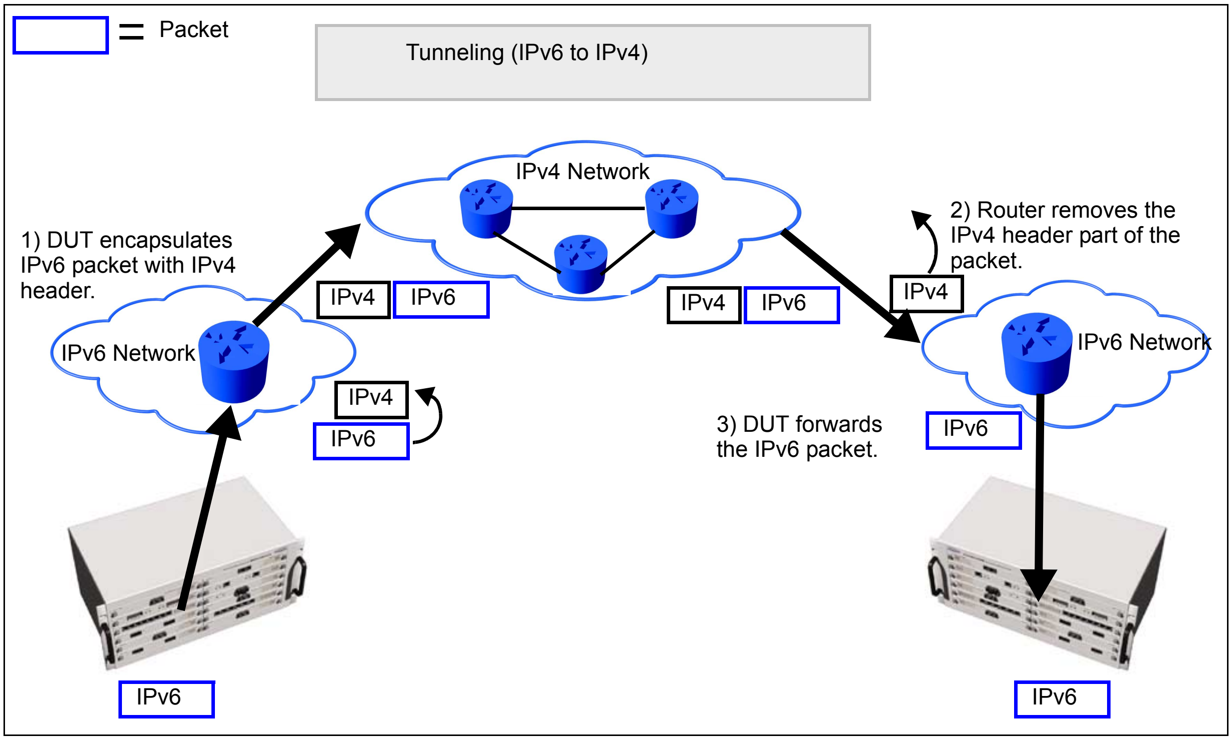

Figure 8-45. Tunneling (IPv6-to-IPv4) Testing

6to4 Tunneling

The *6to4* tunneling method (RFC 3056) uses a well-known prefix of "2002:"which indicates to routers that the packet is encapsulated and sent using an IPv4 tunnel. The syntax is: 2002:v4addr::/48, where v4addr represents the 32-bit IPv4 address of the tunnel destination and is used to encapsulate the IPv4 packet. The IPv4 header has a protocol type of 41 to indicate that the next protocol is an IPv6 header.

The 6to4 tunneling method uses automatic tunneling.

6over4 Tunneling

6over4 (RFC 2893) uses two methods for creating tunnels: automatic or configured. Similar to the 6to4 method, the IPv4 address of the packet is used to encapsulate the IPv6 packet. The IPv4 header has a protocol type of 41 to indicate that the next protocol is an IPv6 header.

To test tunneling:

1. Set up the DUT for the appropriate tunneling mechanism.
2. On the **Cards** tab, check that the **IPv6 Capable** field and the **Addr Resolution** field for both the source and destination ports are checked.
3. On the **Test Setup>Test Iterations** tab, set the **Duration Time (Sec.)** field to at least 60 seconds.
   This value allows the DUT enough time to stabilize its traffic.
4. On the **SmartFlows>IP** tab, specify a destination IP address in the **DST** field that matches the IP address of the tunnel in the DUT.

During the learning phase of the test, if the *Stop on error* field is clear on the *Options>General* tab, an address resolution error is displayed for the SmartBits port on the IPv4 side of the tunnel. Click **Ignore**.

In addition, keep in mind the following tips:

- Be sure to configure the DUT appropriately, so that the DUT can receive and route SmartBits test traffic.
- The flow's destination IP address must match the address of the tunnel that you configured in the DUT. For example, in 6to4 tunneling, start the destination IP address with "2002:" so that all packets use the same tunnel in the DUT.
- Depending on how the destination of the tunnel is configured, prior to the running the test, it may be necessary to populate the DUT IPv4 ARP cache with the IP addresses that will be used on the IPv4 side of the tunnel. For example, if the destination address of the DUT tunnel is the same as the destination address of the destination SmartBits port, ARPing is not necessary.

Network Address Translation (NAT-PT)

The NAT-PT method allows IPv6 hosts to communicate with IPv4 hosts and vice versa by translating IPv4 and IPv6 protocols and addresses.

A NAT device, sometimes known as an Application Level Gateway (ALG), is used that converts the IPv6 addresses to IPv4 addresses (and vice versa) by mapping one to the other. The address assignments can be done statically or dynamically. *Static* address mapping allows you to configure a specific IPv4 address that the DUT should map to a specific IPv6 address. *Dynamic* mapping allows the DUT to draw from a pool of IPv4 addresses to map (dynamically) to an IPv6 address.

SmartFlow supports IPv6-to-IPv4 and IPv4-to-IPv6 address translation. Although other network address protocol translation methods exist, SmartFlow has only been tested with NAT-PT.

Figure 8-46 on page 303 illustrates an example of dynamic address mapping with NAT-PT.

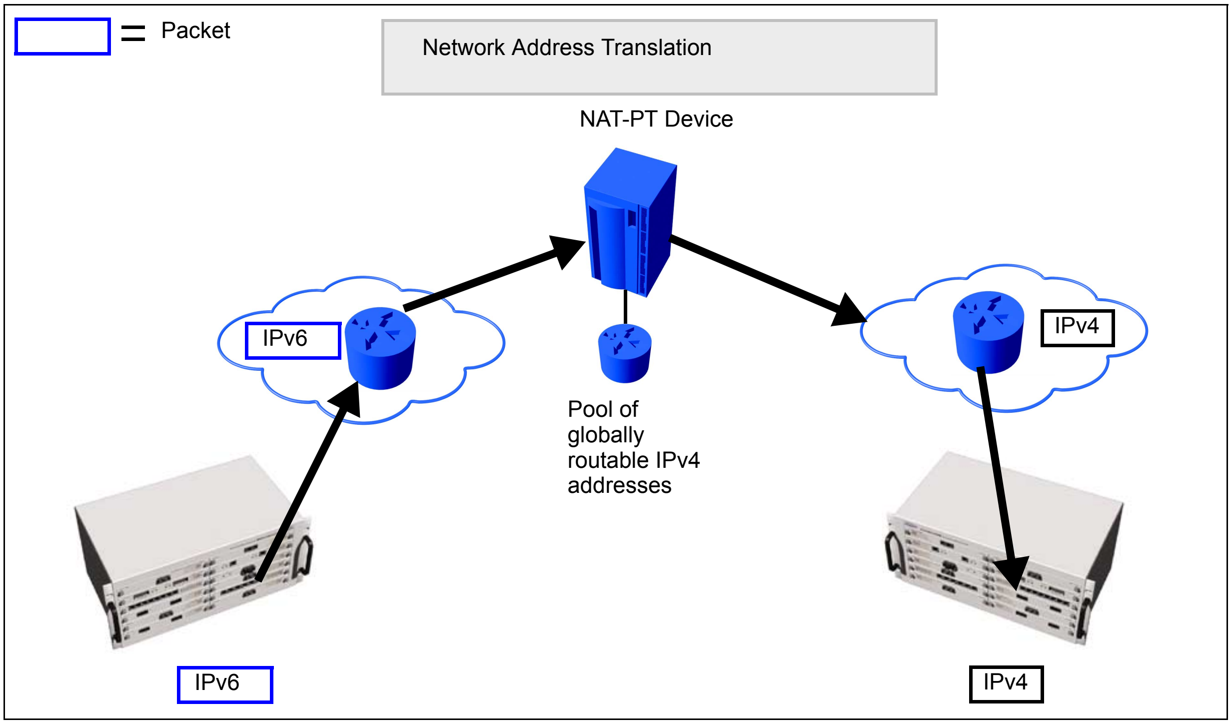

Figure 8-46. Network Address Translation Testing

Follow these tips when using NAT-PT for testing:

- Test the forwarding capability of the NAT-PT device beforehand, so that it does not become a potential bottleneck.
- Consider installing two systems, one for load balancing and one for redundancy, if needed.
- Configure the DUT appropriately, so that the DUT can receive and route SmartBits test traffic.
- Specify on the DUT whether static or dynamic address mapping is being used.
- Configure the mapping table in the DUT to include addresses that you want to use in SmartFlow tests.
- Change the *IP's next protocol* field on the *SmartFlows>Traffic* tab to *TCP* or *UDP* to prevent the possibility of the DUT discarding any packets.

Note: Prior to running the test, it may be necessary to populate the DUT IPv4 ARP cache on the IPv4 side of the tunnel.

QoS Testing

To test QoS in IPv6, set up groups with high, medium, and low flows as explained in *"Sample Flow/Group Setup" on page 130*. Also ensure that IPv6 is enabled for the ports and flows.

Verify that the type of Diff Serv (traffic class) level for the flow (specified on the *IP* tab) is also enabled in the DUT.

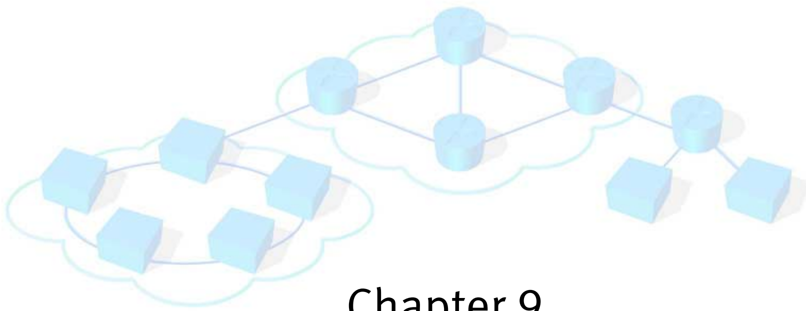

Chapter 9

Throughput Tests

In this chapter...

About the Throughput Tests

The Throughput tests determine the maximum transmission rate at which the DUT can forward IP traffic with no frame loss or at a user-specified acceptable frame loss. By increasing the transmission rate at specified levels, the DUT capacity can be determined. Since prioritization is not visible until there is congestion and thus frame loss, first run a Throughput test to determine the maximum rate without frame loss. Then use the Frame Loss test to identify where frame loss occurs and how priority is handled.

Testing throughput helps answer these questions:

- What is the maximum transmission rate at which the DUT can forward different priority or VLAN levels of IP traffic with no frame loss or at an acceptable frame loss?
- How well does the DUT forward different priority traffic types when traffic is transmitted to the DUT?
- When frame loss occurs, which class of traffic is more likely to be dropped?

Standard and Asymmetric Throughput Testing

SmartFlow offers two types of Throughput tests: standard and asymmetric.

The Standard Throughput test provides all the existing controls contained in previous SmartFlow releases. The option to switch between units (%load/Mbps/Kbps/bps/fps) is not incorporated.

The Asymmetric Throughput test is an additional level of testing that contains extra functionalities. Unlike the standard test, the asymmetric test is able to apply two sets of throughput parameters to two different groups of ports. Throughput is reported asymmetrically - and twice - for two sets of ports that are transmitting bidirectionally.

All of the following information is pertinent to both throughput tests:

- Purpose of tests
- Test methodology
- Determination of next rate
- Search modes (binary, step, and combo)
- Adjustment of test parameters
- Interpretation of test results.

For procedures to enable/set up an Asymmetric or Standard Throughput test, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Asymmetric Port Throughput Test

The Asymmetric Port Throughput test extends the throughput reporting as documented in RFC 2544. Whereas RFC 2544 calls for reporting throughput once and globally for all ports, this new throughput test:

- Reports throughput asymmetrically twice for two sets of ports that transmit bidirectionally: upstream and downstream ports. (A group selection field is available that allows only one group to be selected for upstream and downstream sides, respectively. Remaining groups can be used as background traffic.) This type of reporting is essential for testing broadband access-based networks in which the upstream ports transmit at different rates than the downstream ports. For example, a cable modem test may involve an upstream port transmitting at 3 Mbps and a downstream port transmitting at 128 Kbps. The same premise applies to DSL and PON-based networks in the broadband access market.
- Configures the initial, maximum, and minimum throughput rates in terms of Mbps, Kbps, bps, and fps as well as the previously existing % format. (This reporting feature is supported on all cards.) This is useful when working with Service Level Agreements (SLAs) or WAN circuits that are provisioned in terms of those units.
- Can switch between units.
- Iterates across frame sizes. Iteration is still possible because the units (such as Mbps, Kbps, bps, etc.) can be internally converted into a percentage. The interframe gap for every frame size can be determined.
- Treats each set of ports in one of the following ways:
  - Independently: If there is packet loss on the upstream side, then the rate on the upstream side is reduced per the search algorithm. If there is packet loss on the downstream side, then the rate on the downstream side is reduced per the search algorithm.

  Note: Unless configured otherwise, each set of ports is treated independently.

  - Constant: If there is packet loss on the upstream side, the rate on the upstream side is not reduced. If there is packet loss on the downstream side, the rate on the downstream side is not reduced.
  - Dependently: If there is any frame loss in either the upstream or downstream side, transmission slows in the upstream side.
- Bases the results for each side on the maximum transmit rate configured for that side. The throughput for each side is reported separately as a percentage (0% - 100%) of the maximum rate for that side. Separate graphs for upstream and downstream ports are plotted.
- Reports the asymmetric throughput result as a function of the user-specified maximum rate rather than the port link speed. For example, the maximum rate can be set equal to a provisioned SLA rate. Consequently, the throughput result relates directly to the SLA and not the overall port speed or VC rate.

Note: A Tx port cannot exist in both upstream and downstream groups.

Traffic transmission varies in many ways. For specific examples, refer to the online Help.

Important: For information on configuring proper test loads, see *“Attempting to Go Beyond Load Capabilities of a Port or VCI” on page 308*. When the Asymmetric Port Throughput test is run on multicast flows, the configuration units on the *Test Setup>Individual Tests* tab should be set to %load since acceptable frame loss is measured in terms of % regardless of the *Units* field value.

Throughput results for the asymmetric port throughput test are a percentage of the maximum rate that is specified on the *Test Setup>Individual Tests* tab.

Attempting to Go Beyond Load Capabilities of a Port or VCI

Load capabilities greatly affect both standard and asymmetric throughput testing results.

Standard Throughput Testing with Ethernet

For Ethernet traffic used in the Standard Throughput test, loads may be varied according to the percentage of maximum port bandwidth. Therefore, ports of 10 Mbps transmit at 10 Mbps when set to 100%. A port of 100 Mbps transmits at 100 Mbps when set to 100%. Each port transmits based on a percentage of its link speed. For Ethernet ports, the Standard Throughput results are reported as a function of the port link speed.

Standard Throughput Testing with ATM and WAN

ATM and WAN behave in a similar manner in regards to Standard Throughput testing. (The following text references ATM channels, but the same considerations apply to running throughput tests over WAN.)

The virtual circuits defined on the ATM port may fill the available bandwidth, but more normally will represent only a small percentage of the port speed. When Standard Throughput tests are run, each virtual circuit is throttled according to the percentage load. For example, a 64 Kbps circuit defined on a port that is to transmit at 50% only transmits at a rate of 32 Kbps. The same virtual circuit running at 75% load operates at 48 Kbps. For ATM ports, the Standard Throughput results are reported as a function of the configured VC rate.

Asymmetric Throughput Testing with Ethernet

Asymmetric Throughput testing must obey the same limits of port speed as defined for Standard Throughput testing. However, with Asymmetric Throughput tests, the load is varied on a port-wide basis.

For the Asymmetric Throughput test, loads may be defined in % utilization, Kbps, Mbps, or bps. Each port transmits at a given load according to the throughput search algorithm. The transmit load for a port will not exceed the link speed of the port, nor the maximum rate set in the asymmetric throughput section on the *Test Setup>Individual Tests* tab. When requested to transmit at 100 Mbps, a 10 Mbps port can only transmit at its limit of 10 Mbps. However, a 100 Mbps port can transmit at 100 Mbps. If the maximum rate is set to 1 Mbps, then the port does not transmit more than 1 Mbps.

For Ethernet ports, the Asymmetric Throughput test results are reported as a function of the maximum rate set in the asymmetric throughput section on the *Test Setup>Individual Tests* tab.

Using ATM with Asymmetric Throughput Testing

An ATM virtual circuit may be considered as a lower bandwidth port where its maximum throughput is limited by the individual circuit configuration. Multiple virtual circuits may comprise a single physical ATM port. When the ATM load is expressed in Mbps/Kbps/

bps, the individual virtual circuit transmission rate may not exceed its configured value. For a 64 Kbps virtual circuit configured on an OC3C ATM port of 155,519 Kbps, the transmission rate never exceeds 64 Kbps on that virtual circuit and only drops below 64 Kbps when the port load itself drops below 64 Kbps. In a similar manner, should two virtual circuits of 64 Kbps each be configured on the same ATM port, their transmission does not drop below 64 Kbps each until the port transmission rate falls below 128 Kbps.

A 100% transmission rate is achieved on each virtual circuit as soon as the port transmission rate meets or exceeds the combined total.

Should the asymmetric throughput loads be defined as a percentage, the same rules apply. However, the effective port load in Kbps will be calculated from the ATM OC3/OC12 setting. The effective load in Kbps will be derived from the port bandwidth (OC3/OC12) multiplied by the percentage load, i.e., 50% of OC3 (155,519 Kbps) gives a port transmission rate of 777,595 Kbps. In this scenario, a single 64 Kbps virtual circuit would not fall below 100% transmission until the port bandwidth falls below 64 Kbps or 0.0411%.

For ATM ports, the Asymmetric Throughput test results are reported as a function of the maximum rate set in the asymmetric throughput section on the *Test Setup>Individual Tests* tab.

Asymmetric Throughput Testing with ATM and Ethernet

When testing a combination of ATM and Ethernet traffic, significant differences in bandwidth may exist between the two interface types. Therefore, throughput should be measured on each interface independently of the other, either by selecting the *Upstream then Downstream* option or choosing the *Independent* field value for the *Backoff basis* mode on the *Test Setup>Individual Tests* tab.

In addition, take into consideration the port utilization of test ports when configuring test loads. If port utilizations are below configured test loads, test ports cannot transmit beyond rates specified in the utilizations on the *Cards* tab.

Asymmetric Throughput Loads That are Defined as Frame Rates

In Asymmetric Throughput testing, the load may also be set in terms of a frame rate. The effective load in bps is calculated internally, by taking the defined fps value and multiplying that by each frame length. Hence the initial load, minimum load, and maximum load may all be generated from user-defined fps values.

Setting Traffic Load per Flow, Using the Custom Loads per Flow Table

Where custom loads per flow have been defined on the *Test Setup>Test Iterations* tab, their effect is to further reduce the effective throughput. The overall port load is further throttled by the combined flows custom loading. For instance, two flows sourced from a port of 100 Mbps have custom loads set to 50% and 25%, respectively. Their combined loading is 75%, and so even when the port is configured to run at 100%, the effective bandwidth used is only approximately 75 Mbps. (Note that this does not account for the effects of preamble. Custom loads per flow do not work with fps test loads.)

Examples of ATM and Throughput

Case 1

ATM - Standard Throughput (throughput loads in percentages; percentage is the only option in standard throughput)

VC1 - 400 Kbps

VC2 - 800 Kbps

The initial percentage load is compared to each VC rate specified on the ATM tab. Each virtual circuit transmits as a percentage of the configured VC rate. When the test load is 50%, then VC1 transmits at a rate of 200 Kbps while VC2 transmits at 400 Kbps. When the test load is 75%, VC1 transmits at a rate of 300 Kbps while VC2 transmits at 600 Kbps.

The total port transmit load is increased or decreased based on each VC rate and the percentage load.

Throughput is reported as a function of the configured VC rate.

Note: The Asymmetric Throughput test provides for setting the transmit loads in terms of percentage utilization, Mbps, Kbps, fps, or bps.

Case 2

ATM - Asymmetric Throughput (throughput loads in percentages)

VC1 - 400 Kbps

VC2 - 800 Kbps

The initial test load is compared to the port speed. When the percentage load is 50% and the port speed is 155,519 Kbps (OC3), the port attempts to transmit 777,595 Kbps. Based on its configured ATM rate, each virtual circuit transmits its proportionate spare of this port load. VC1 transmits 400 Kbps and VC2 transmits 800 Kbps since the combined rate is less than 777,595 Kbps. The two VC actual rates are reduced from the configured rates if:

- There is packet loss at 1,200 Kbps (or less) for the port.
- The initial, step, or maximum load is configured to less than 1,200 Kbps.

The port transmit load is increased or decreased based on the percentage load, the port speed, and the VC rate.

The throughput result is reported as a function of the throughput maximum rate set on the *Test Setup>Individual Tests* tab.

Note: If you want the search algorithm and throughput result to be based on the total of all provisioned virtual circuits, then set the maximum load to 0.7716% (1,200 Kbps/ 155,519 Kbps).

Case 3

ATM - Asymmetric Throughput (with throughput loads in Mbps, Kbps, fps, or bps)

VC1 - 400 bps

VC2 - 800 bps

The initial load (specified in Mbps, Kbps, fps, or bps) is compared to the port speed and its effective transmit rate. For example, if the initial load is 200,000 Kbps, then the port attempts to transmit 200,000 Kbps. Based on its configured ATM rate, each virtual circuit transmits its proportionate share of this port load. VC1 transmits 400 Kbps and VC2 transmits at 800 Kbps since the combined rate is less than 200,000 Kbps. The two VC rates are reduced from the configured rates if:

- There is packet loss at 1,200 Kbps (or less) for the port.
- The initial, step, or maximum load is configured to less than 1,200 Kbps.

The port transmit load is increased or decreased based on the percentage load, the port speed, and the VC rate.

The throughput result is reported as a function of the throughput maximum rate set on the *Test Setup>Individual Tests* tab.

Note: If you want the search algorithm and throughput result to be based on the total of all provisioned virtual circuits, then set the maximum load to 1,200 Kbps.

In this third case, if two VCs are configured (one at 400 Kbps and the other at 800 Kbps), you may want to set the following for the throughput loads:

Initial rate = 100 Kbps

Minimum rate = 100 Kbps

Maximum rate = 1,200 Kbps

Test Methodology

Throughput tests the maximum transmission rate at which the DUT can forward traffic for existing flows and/or groups without frame loss. SmartFlow calculates frame loss as:

Frame Loss = Number of Frames Transmitted - Number of Frames Received

The general methodology of either Throughput test is as follows:

- The test finds the maximum transmission rate at which the DUT can forward traffic for existing flows and/or groups.
- During the test, the transmission rate changes within a user-specified range, depending on whether the current rate passed or failed. If a single frame is dropped or if loss exceeds the specified acceptable loss rate, the test fails and it stops or is repeated at a lower load, depending on the search mode selected. If you are specifying rates per flow, any flow that fails causes a failure for that test iteration.

- When the test finishes depends on the search mode:
  In Binary search mode: The test finishes when the difference between the rate of the last passed test and last failed test is less than or equal to the percent specified in the *Resolution (%)* field on the *Test Setup>Individual Tests* tab. For more information, see *"When Does a Test End in Binary Mode?" on page 314*.
  In Step search mode: The test finishes when it reaches the specified maximum load or it fails. (Throughput is determined based on the last passing rate.)

Note: During test setup, you can specify a percentage of frame loss that is acceptable and for which the test will not fail in the *Acceptable frame loss (%)* field on the *Test Setup>Individual Tests* tab. However, if a non-zero value is entered in this field, the test will not conform to the RFCs (1242 and 2544) specification of throughput. For multicast flows, acceptable frame loss is measured in terms of % regardless of the *Units* field value.

The number of frames sent is determined by the length of time the test is to run, frame size, and load. SmartFlow only counts frames generated by the sending SmartBits port, not frames sent by the DUT. Keepalive and routing update frames are not counted as received frames.

How SmartFlow Determines Throughput

SmartFlow determines whether a test passes or fails by comparing the percentage of frame loss to the value specified in the *Acceptable frame loss (%)* field on the *Test Setup>Individual Tests* tab.

Table 9-1. Frame Loss

If Frame Loss % is...	Then ...
Less than the value in the *Acceptable frame loss (%)* field	The test passes. The current rate becomes the throughput rate.
Greater than the value in the *Acceptable frame loss (%)* field	The test fails. The last passing rate becomes the throughput rate.

How SmartFlow Determines the Next Throughput Rate

If the initial rate is 100%, the first frame rate for a given frame length is the maximum rate for the topology and speed of the transmitting SmartBits port. If all frames from the transmitting SmartBits port(s) are received by the receiving port(s), SmartFlow attempts no further trials and the maximum frame rate is recorded as the throughput rate.

If the first trial fails (loss exceeds the specified acceptable loss rate), the search mode chosen in the test setup on the *Test Setup>Individual Tests* tab determines how SmartFlow calculates the rate for the second and subsequent trials:

- Binary search mode uses a binary (scaled) search to determine a rate between the last failed rate and the last successful rate.

- Step search mode increases the rate by the same percentage (step) added to the current rate.
- Combo search mode starts with the step search mode, stepping up the rate until the test fails. Then it performs binary searches.

For more information, see *"Search Modes"*.

Search Modes

A Throughput test progresses differently depending on the selected mode (binary, step, or combo) in the test setup on the *Test Setup>Individual Tests* tab. Select a mode based on the type of information that you want from the test. This table lists the type of information available in each mode.

If you want to know...	Select...
The maximum transmission rate for traffic.	Binary
The relative frame loss at different transmission rates for traffic.	Step
Both maximum transmission rate and relative frame loss.	Combo

Note: For non-ATM cards, load variation is only applied to port loads, not flow loads. Custom loads per flow are still dependent upon the changes in port load. Basically, flows using custom loads will not be able to have loads varied in the same manner as ATM flows.

See *"Binary Search Mode"* and *"Step Search Mode" on page 317* for more information about each mode.

Binary Search Mode

The binary search mode determines the maximum transmission rate for traffic but not relative frame loss at different transmission rates. In binary search mode, the test starts at the transmission rate specified in the *Initial rate (%)* field on the *Test Setup>Individual Tests* tab and increments or decrements based on the specified backoff ratio. For the first iteration, the current rate is the initial rate.

The transmission rate of traffic is increased or decreased based on whether the DUT passes or fails.

- DUT passes at the current transmission rate. This means frame loss is less than or equal to the value specified in the *Acceptable frame loss (%)* field.
  If the DUT passes, the transmission rate is increased using the following equation:

New Rate = Current Rate + [(Backoff Rate) x
(Last Failed Rate % - Current Rate)]

(If the DUT has never failed prior to the current iteration, SmartFlow uses the maximum rate instead of the last failed rate.)

- DUT fails at the current transmission rate. This means frame loss is greater than or equal to the value specified in the *Acceptable frame loss (%)* field.

  If the DUT fails, the next transmission rate is reduced using the following equation:

  New Rate = Last Passing Rate + [(Backoff Rate) x
  (Current Failed Rate - Last Passing Rate %)]

  If the DUT has never passed prior to the current iteration, SmartFlow uses the minimum rate instead of the last passing rate.

After the transmission rate is adjusted up or down accordingly, the test continues at the new, adjusted transmission rate. After each pass or failure, SmartFlow adjusts the transmission rate up or down accordingly.

When Does a Test End in Binary Mode?

In binary mode, each test iteration at a particular rate continues until one of the following occurs:

- The value specified in the *Maximum rate (%)* field on the *Test Setup>Individual Tests* tab is reached.

  The *Continue beyond min or max* checkbox can affect when a test ends in binary mode. When this checkbox is selected, the *Maximum rate (%)* and *Minimum rate (%)* fields serve to focus on a range and thus speed up testing, rather than determine the absolute limits of the test. If SmartFlow cannot determine throughput within this range, it expands the range to either 0% or 100% (whichever applies) and continues searching for throughput. This option eliminates any guesswork in attempting to define a range that will successfully yield throughput. See *“Example 3: Throughput is Higher than Maximum Rate Specified” on page 316* for an example.

- The value specified in the *Minimum rate (%)* field on the *Test Setup>Individual Tests* tab is reached.

  When the *Continue beyond min or max* checkbox is selected, if the test fails at the minimum rate specified, SmartFlow expands the binary search range to between the value in the *Minimum rate (%)* field and 0%. The test continues until the throughput rate is reached or the rate is 0%. See *“Example 4: Throughput is Lower than Minimum Rate Specified” on page 317* for an example.

- The throughput measurement is reached within the specified minimum/maximum range.

  The throughput measurement is reached when the required adjustment percentage (up or down) in the transmission rate is smaller than the specified *Resolution (%)* field value on the *Test Setup>Individual Tests* tab. Once this occurs, the test stops.

For example, if the new, adjusted transmission rate is calculated to be 5.00% higher or lower than the current rate, the adjustment is made only if the *Resolution (%)* field value is less than 5.00%. The smaller the resolution, the more accurate the results. Refer to the online Help for a definition of the *Resolution (%)* field.

As shown in *Figure 9-1*, when the test finishes also depends in part on whether the *Continue beyond min or max* checkbox is selected.

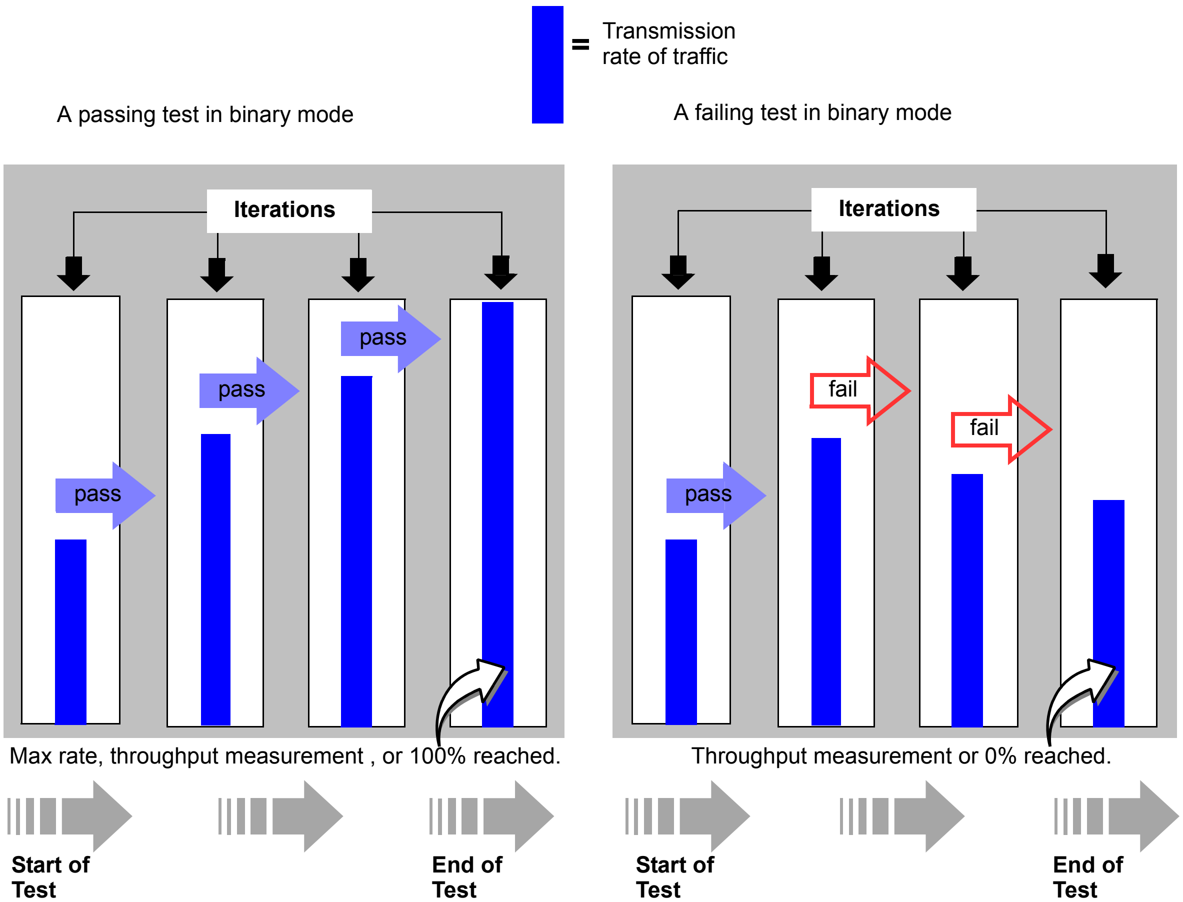

Figure 9-1. Throughput Test Using Binary Search mode

Binary Search Mode Examples

Example 1: Backoff of 50%

Let us say you select the following values on the *Test Setup>Individual Tests* tab:

Initial rate (%):	50.00
Maximum rate (%):	100.00
Backoff (%):	50.00
Minimum rate (%):	0
Resolution (%):	15.00

According to this test setup, the test runs as follows:

1. The test passes at 50%.
2. SmartFlow increases the rate to 75% and the test fails.
3. SmartFlow drops the rate to 62.5%, but since the resolution is 15%, the test stops at this point because the difference between 75 and 62.5% is less than the resolution amount.

Example 2: Backoff of 7 0.71% ($1/(\sqrt{2})$)

Let us say you select the following values on the *Test Setup>Individual Tests* tab:

Initial rate (%):	100.00
Maximum rate (%):	100.00
Backoff (%):	70.71
Minimum rate (%):	0
Resolution (%):	1.00

According to statistical theory, by using 70.71% ($1/(\sqrt{2})$) for the backoff rate, it allows SmartFlow to arrive at the throughput rate faster than by using another percent.

According to this test setup, the test runs as follows:

1. The test fails at 100%.
2. SmartFlow compares the rate between 100% and 0% (considered the last passing rate) and drops the rate to 70.71% of the difference.

Example 3: Throughput is Higher than Maximum Rate Specified

Let us say you select the *Continue beyond min or max* checkbox and select the following values on the *Test Setup>Individual Tests* tab:

Initial rate (%):	40.00
Maximum rate (%):	80.00
Backoff (%):	50.00
Minimum rate (%):	0
Resolution (%):	15.00

According to this test setup, the test runs as follows:

1. The test starts at 40% and passes each iteration until it reaches 80%.

2 SmartFlow extends the maximum range and does a binary search between 80% and 100%.

3 If the test still passes at 90%, SmartFlow continues to do binary searches until throughput is found or the rate reaches 100%, at which point the test ends.

Example 4: Throughput is Lower than Minimum Rate Specified

Let us say you select the *Continue beyond min or max* checkbox and select the following values on the *Test Setup>Individual Tests* tab:

Initial rate (%):	50.00
Maximum rate (%):	80.00
Backoff (%):	50.00
Minimum rate (%):	20
Resolution (%):	15.00

According to this test setup, the test runs as follows:

1 The test starts at 50% and fails each iteration until it reaches 20%.

2 SmartFlow extends the minimum range and does a binary search between 20% and 0%.

3 If the test still fails at 10%, SmartFlow continues to do binary searches until throughput is determined or the rate drops to 0%, at which point the test ends.

Step Search Mode

The step search mode provides relative frame loss at different transmission rates for traffic. The *Continue beyond min or max* checkbox on the *Test Setup>Individual Tests* tab can affect when a test ends in step search mode. By selecting this checkbox, the *Maximum rate (%)* and *Minimum rate (%)* fields help speed up testing by focusing on a range, instead of determining the limits of the test. When selected, if SmartFlow cannot determine throughput within this range, it expands the maximum rate to 100% and continues searching for throughput.

The test starts at the transmission rate that is specified for the *Initial rate (%)* test parameter. During each trial, the transmission rate increases according to the *Step rate (%)* specified.

The test continues until either the current rate reaches the specified *Maximum rate (%)* field value or the DUT fails in a specific iteration, at which point the last passing rate becomes the throughput.

Notes:
- In step search mode, the *Backoff (%)* and *Resolution (%)* parameters are not used.
- Load starts from the initial value specified and increments at steps whose values are also specified until the maximum load is reached. If the current load is less than the maximum load, but the difference between the two is

larger than the acceptable loss and also larger than the step increment, the current load is set to the maximum load.

If the *Continue beyond min or max* checkbox was selected and the test passes at the maximum rate specified, SmartFlow expands the binary search range to between the value in the *Maximum rate (%)* field and 100%. The test continues until SmartFlow can determine throughput or the throughput is 100%.

During the test, the Results window appears, allowing you to view the test as it progresses. Once the test has completed, you can save the results to `.csv` files.

Figure 9-2 illustrates how a Throughput test using the step search mode works.

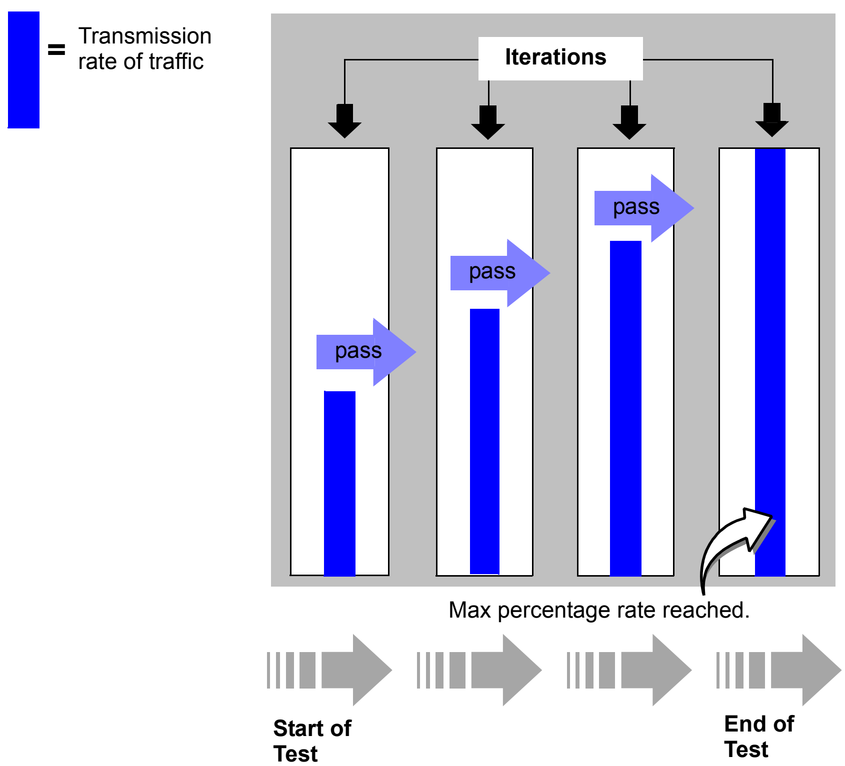

Figure 9-2. Throughput Test Using Step Search Mode

Combo Search Mode

The combo search mode is a combination of the binary and step search modes. This method may search throughput results faster than the other methods. The combo search mode works as follows:

1. SmartFlow increases the current passing rate by the percentage specified in the *Step rate (%)* field until the test fails.

2. Once the test fails, SmartFlow does a binary search for a rate between the last passing rate and the current failing rate.

   When the *Continue beyond min or max* checkbox is selected and the test fails at the maximum rate specified, SmartFlow expands the binary search range to between the value in the *Maximum rate (%)* field and 100%. The test continues until the throughput rate is reached or the rate is 100%. See *"Example 3: Throughput is Higher than Maximum Rate Specified" on page 316* for an example.

3. SmartFlow continues to do binary searches until it finds the maximum throughput rate.

For more information about the binary and step search modes, see sections *"Binary Search Mode" on page 313* and *"Step Search Mode" on page 317*.

Note: If the *Step rate (%)* field value is smaller than or equal to the *Resolution (%)* field value, the combo search mode behaves exactly the same as step search mode.

Figure 9-3 illustrates how a Throughput test works when using the combo search mode.

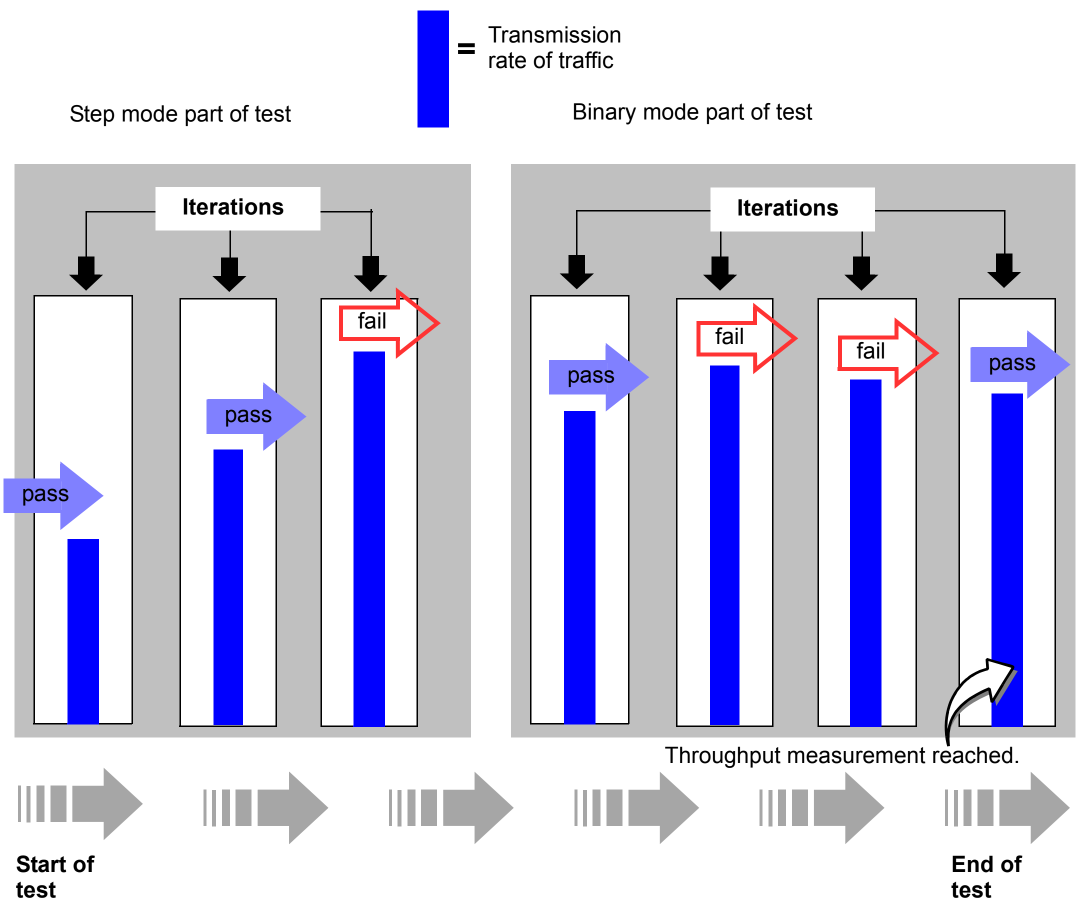

Figure 9-3. Throughput Test Using Combo Search Mode

Example:

Let us say you select the following values on the *Test Setup>Individual Tests* tab:

Initial rate (%):	40.00
Maximum rate (%):	100.00
Backoff (%):	50.00
Minimum rate (%):	0
Resolution (%):	15.00
Step rate (%):	10.00

According to this test setup, the test runs as follows:

1. The test performs the step algorithm starting from 40% until it reaches 80% and then fails.
2. The algorithm now enters the binary search mode and performs a search for a rate between 70 (last passing rate) and 80 (current failing rate).

Adjusting the Throughput Test Parameters

This table provides ways to rerun a Throughput test with different parameter values based on the test results.

If test results show...	Try one or more of the following...
Frame loss occurring at or near the start of the test (In this case, the actual throughput of the DUT may be lower than any levels run during the test. Rerun the test.)	• Run the test again with a lower *Initial rate (%)* setting. • Run the test again with a lower *Step rate (%)* setting.
No frame loss occurs during the test (In this case, the actual throughput of the DUT may be higher than any levels run during the test. Rerun the test.)	• Run the test again with a higher *Initial rate (%)* setting. • Run the test again with a higher *Maximum rate (%)* setting.

Interpreting Throughput Test Results

Results of a Throughput test automatically appear in the main window at the conclusion of the test. You can view test results for this test in graphical or tabular form. In tabular form, you can view a summary, a detailed report, or a specialized (such as stray frames) report.

Notice in the test results that where frame loss starts occurring, the DUT also starts prioritizing how it handles groups and/or flows that it forwards.

Note: The format of results can vary depending on whether or not you used the *Frame size automation (all flows, with CRC)* option on the *Test Setup>Test Iterations* tab to vary frame sizes option during the test. For more information, see *"Using Frame Size Automation (Global for All Flows)" on page 236*.

Chart and Summary Throughput Unit Options

Some DUTs have interfaces with bandwidth that is less than the rate at which the SmartBits card/module is transmitting. Some DUTs have interfaces that only support a particular bandwidth. In order to be able to compare the receive rate of the DUT to the rate being used in the test, you can specify the units by which to view the throughput in chart or summary results.

You can view throughput in chart and summary results using any of these units:

- %Throughput
- Frames/Sec
- Bits/Sec (Without L2header & CRC).

In addition, if the test was run using frame size automation, you can select the *Per Frame Size* option to view results by frame size instead of test load iteration.

The Frames /Sec and Bits/Sec (Without L2header & CRC) units refer to the receiving rate and are presented in scientific notation if the value goes beyond six digits. (See *Figure 9-6* for an example.) For information about scientific notation, see *"Scientific Notation in Chart Results" on page 81*.

To specify the unit of throughput for chart or summary results:

1. Right-click anywhere on the **Chart** or **Summary** results tab.
2. Click the appropriate unit(s).

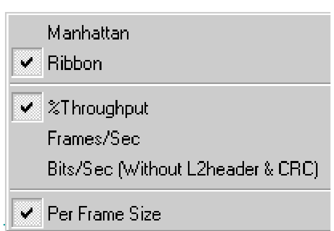

For examples of how chart results appear with each of these options, see *"Chart Results (Throughput)" on page 322*.

For a description of the Throughput chart and summary results, right-click menu format, and unit options, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Chart Results (Throughput)

Throughput chart results vary according to the unit that you use to display throughput, and whether or not you are using the *Frame size automation (all flows, with CRC)* option on the *Test Setup>Test Iterations* tab to vary frames.

Throughput chart results for a test run *without* varying frame sizes illustrate a pattern of frame loss across various loads for each group and/or flow, as shown in *Figure 9-4*.

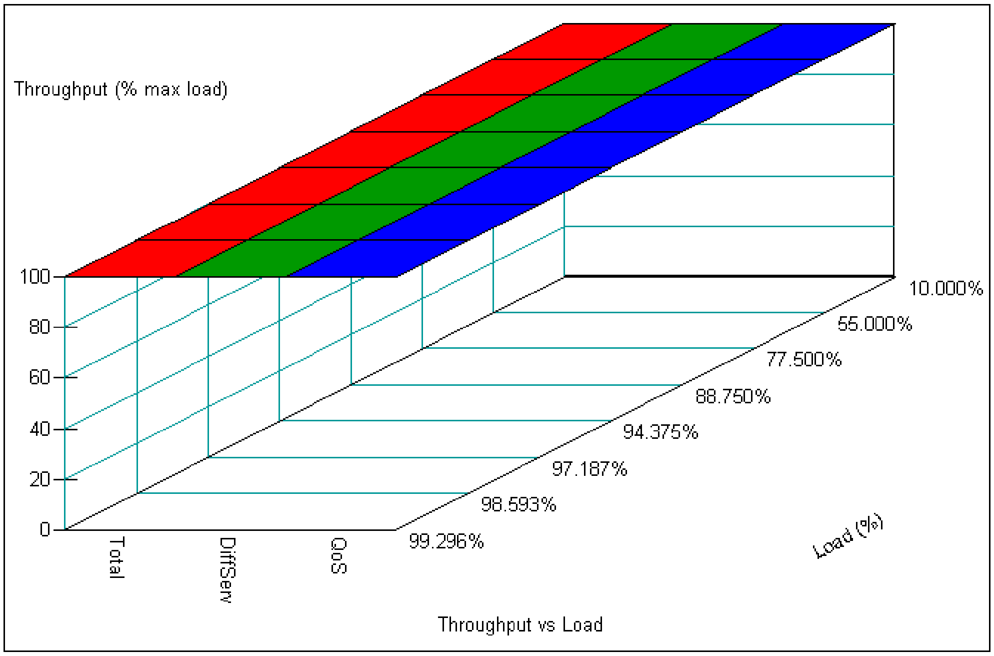

Figure 9-4. Throughput Chart Results without Varying Frame Sizes

You can change the unit by which to view results by right-clicking on the chart. See *"Chart and Summary Throughput Unit Options" on page 321*.

For a description of each field on the Throughput *Chart* results tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Throughput Chart Results with Frame Size Automation

Throughput chart results for a test run with frame size automation show throughput across various loads for each test iteration, as shown in *Figure 9-5*.

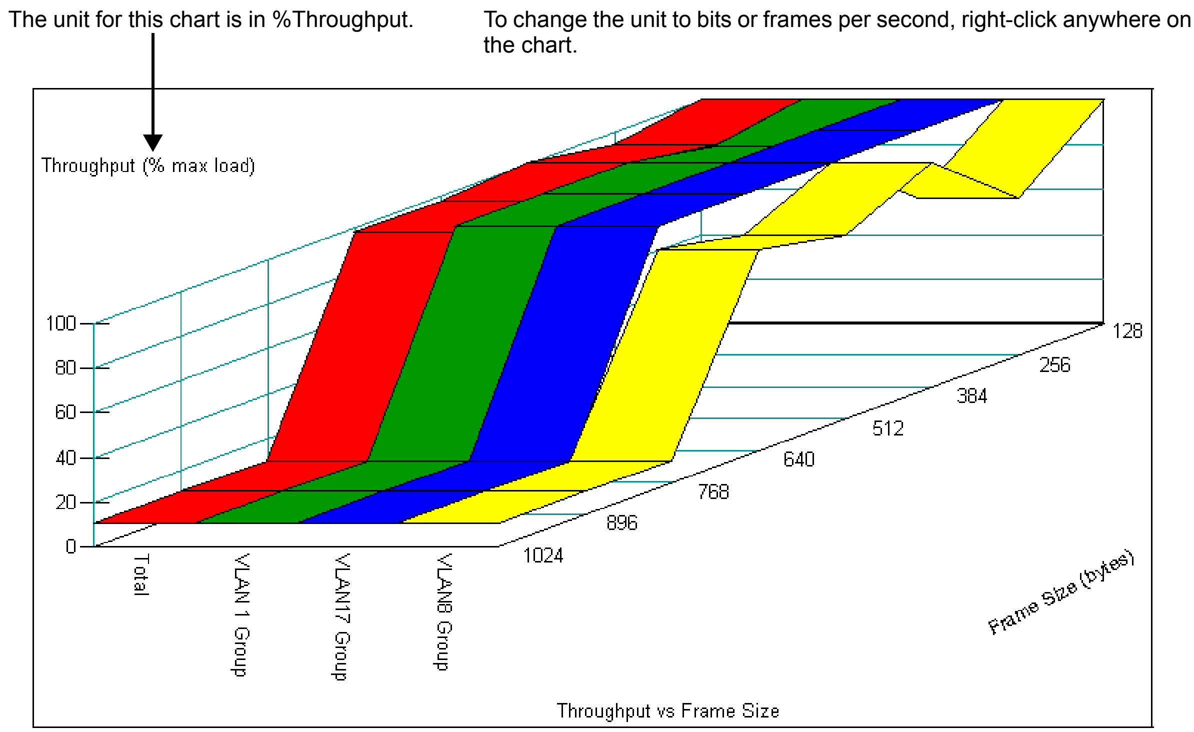

Figure 9-5. Throughput Chart Results with Varying Frame Sizes in %Throughput Units

Notice that the unit for the chart in *Figure 9-5* is *%Throughput.* Flows in two of the groups are transmitted from 100 Mbps Ethernet transmitting ports and one from a Gigabit Ethernet port. However, each of the ribbons representing the groups is roughly the same in terms of throughput percentage.

For a description of each field on the Throughput *Chart* results tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Now compare the test results that appear in *Figure 9-5* in *%Throughput* format to the same results in *Bits/Sec (Without L2header & CRC)*, as shown in *Figure 9-6 on page 324*. Notice the large difference between the number of bits transmitted from 100 Mbps ports and the bits transmitted from Gigabit ports, whereas if the format was *%Throughput,* the difference is not visible.

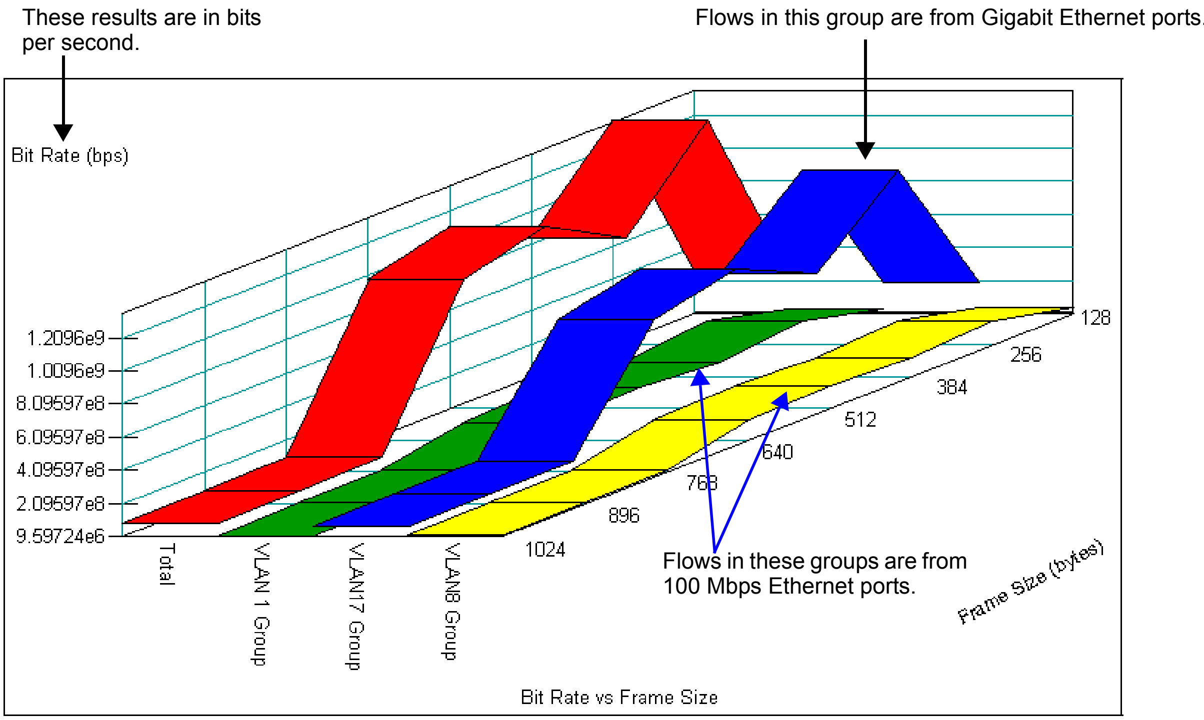

Figure 9-6. Throughput Chart Results with Varying Frame Sizes in Bits per Second Units

The number of bits in the example in *Figure 9-6* is shown in scientific notation. For more information about scientific notation, see *"Scientific Notation in Chart Results" on page 81*.

The same test results that are in bits per second units in *Figure 9-6* can additionally be presented per frame size instead of per test load iteration. *Figure 9-7 on page 325* shows Throughput in bits per second for each frame size in the test.

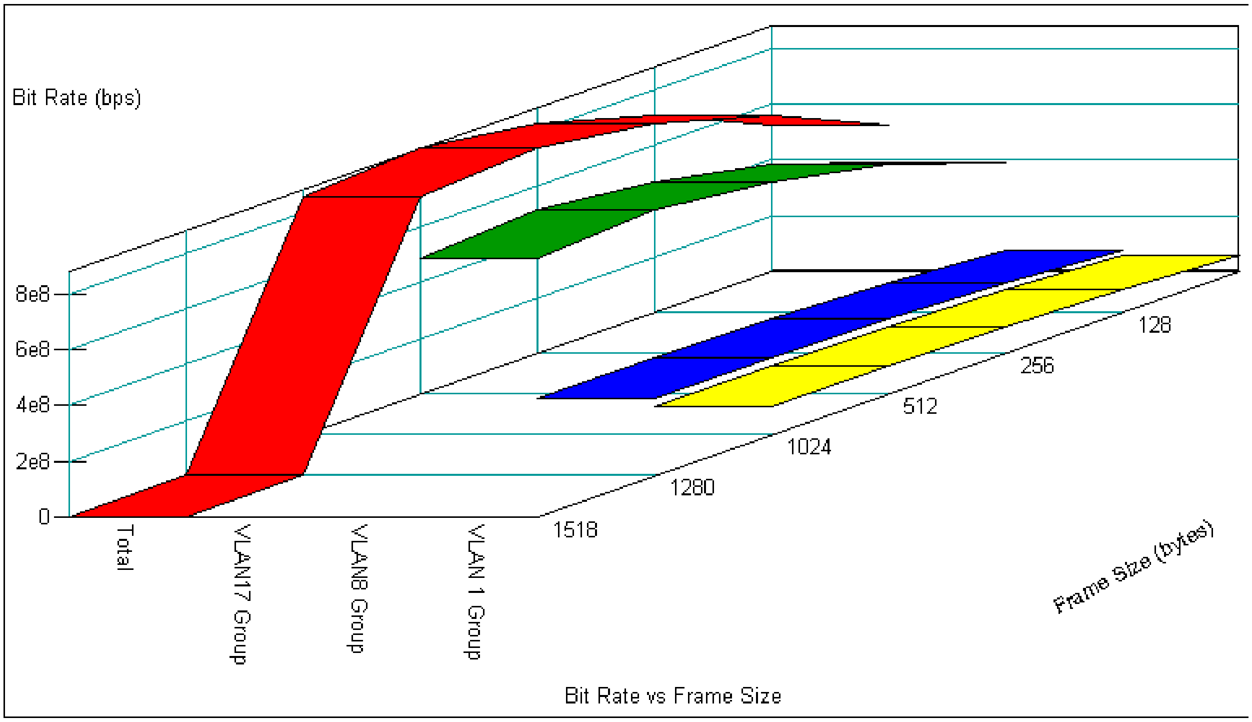

Figure 9-7. Throughput Chart Results with Varying Frame Sizes in Bits per Second Units per Frame Size

Summary Results Table (Throughput)

The Throughput summary results without frame size automation show the percentage of frames lost at the specified load intervals for flows and groups involved in the test.

Name/Load	56.00000	57.00000	58.00000	59.00000	60.00000	61.00000	62.00000
Total					14.2692		
A Group					14.2692		

For a description of each field on the Throughput *Summary* results tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

The Throughput summary results with frame size automation show aggregate frame loss that occurred for all groups and/or flows at each frame size. Even if the *Acceptable frame loss* field on the *Test Setup>Individual Tests* tab is set to 0, the percent of frame loss displays as zero as long as SmartFlow is able to determine throughput. Frame loss appears as non-zero if the passing throughput rate is less than the maximum rate specified for the test.

FrameSize	128	256	384	512
Throughput (% max load)	100.0000	98.5938	7.1875	9.4375
Frame Loss (%)	0.0000	0.0000	0.0000	0.0000

Detailed Results Table (Throughput)

For flows and groups involved in the test, the Throughput *Detail* results tab shows the transmission rate, number of frames sent, received, lost, and percent of loss. Results display total throughput at the start of each new test iteration. Final throughput is determined at the end of the test.

If the percentage in the *Load* column increases steadily from one iteration to the next until there is frame loss, the test used the step search mode to increase the new rate. If the load increases, then decreases, then increases again, the test used the binary search mode to arrive at the next rate. For more information about how these modes work or how SmartFlow determines the next rate, see *"Search Modes" on page 313*.

When the number of frames sent is equal to the number received, there is no frame loss.

Notes:

- Detailed Throughput test results may show "empty" iterations below the point at which results for a successful test completion are displayed. SmartFlow inserts these iterations as placeholders at the start of the test since at that point the number of iterations needed is unknown.
- Since throughput is measured in terms of port load, not individual flows, N/A is displayed in the *Throughput* column for flows.
- When reporting the %load, Smartflow includes both IFG and preamble in the calculation. This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

Figure 9-8 shows results for a test run with frame size automation:

The next frame size begins here.

This is the throughput for this iteration.

This is the current load for the flow.

This test was run using binary mode, so the load is raised in an attempt to determine throughput.

Number of frames sent minus received equals frames lost.

Name	Time	FrameSize	ILoad	TxFrames	RxFrames	LostFrames	Lost (%)	Throughput
D 2-2->2-1	11/16/01 2:13:59 PM	64	16.60100	247044	247044	0	0.00	N/A
Total	11/16/01 2:14:26 PM	256	10.00000	317016	317016	0	0.00	10.00
Vlan1	11/16/01 2:14:26 PM	256	10.00000	90576	90576	0	0.00	10.00
Vlan 2	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	10.00
Vlan 3	11/16/01 2:14:26 PM	256	10.00000	135864	135864	0	0.00	10.00
Vlan 4	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	10.00
A 1-1->1-4	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	N/A
A 1-2->1-4	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	N/A
B 1-4->1-1	11/16/01 2:14:26 PM	256	5.00000	22644	22644	0	0.00	N/A
B 1-4->1-2	11/16/01 2:14:26 PM	256	5.00000	22644	22644	0	0.00	N/A
C 1-5->2-2	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	N/A
C 1-6->2-2	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	N/A
C 2-1->2-2	11/16/01 2:14:26 PM	256	10.00000	45288	45288	0	0.00	N/A
D 2-2->1-6	11/16/01 2:14:26 PM	256	5.00000	22644	22644	0	0.00	N/A
D 2-2->2-1	11/16/01 2:14:26 PM	256	5.00000	22644	22644	0	0.00	N/A
Total	11/16/01 2:14:52 PM	256	55.00000	1743504	1403184	340320	19.52	10.00
Vlan1	11/16/01 2:14:52 PM	256	55.00000	498144	452516	45628	9.16	10.00
Vlan 2	11/16/01 2:14:52 PM	256	55.00000	249072	249072	0	0.00	10.00
Vlan 3	11/16/01 2:14:52 PM	256	55.00000	747216	452524	294692	39.44	10.00
Vlan 4	11/16/01 2:14:52 PM	256	55.00000	249072	249072	0	0.00	10.00
A 1-1->1-4	11/16/01 2:14:52 PM	256	55.00000	249072	219909	29163	11.71	N/A
A 1-2->1-4	11/16/01 2:14:52 PM	256	55.00000	249072	232607	16465	6.61	N/A
B 1-4->1-1	11/16/01 2:14:52 PM	256	27.50000	124536	124536	0	0.00	N/A
B 1-4->1-2	11/16/01 2:14:52 PM	256	27.50000	124536	124536	0	0.00	N/A
C 1-5->2-2	11/16/01 2:14:52 PM	256	55.00000	249072	152435	96637	38.80	N/A
C 1-6->2-2	11/16/01 2:14:52 PM	256	55.00000	249072	163384	85688	34.40	N/A
C 2-1->2-2	11/16/01 2:14:52 PM	256	55.00000	249072	136705	112367	45.11	N/A
D 2-2->1-6	11/16/01 2:14:52 PM	256	27.50000	124536	124536	0	0.00	N/A
D 2-2->2-1	11/16/01 2:14:52 PM	256	27.50000	124536	124536	0	0.00	N/A

Figure 9-8. Throughput Detail Results with Frame Size Automation

The *Name* column is fixed, so that when you scroll to the right it is always displayed.

Note: If you get an error message that states "Received count is larger than transmitted count (Duplicate Packets might have been received)," clicking **Ignore** may give corrupt results in the *LostFrames* and *Lost (%)* columns for those flows experiencing the duplicate frames.

For a description of each field on the Throughput *Detail* results tab as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Stray Frames Results Table (Throughput)

The Throughput *Stray Frames* results tab shows the frames per port that were received but were not destined (expected) for that port. It also lists each flow (not SmartFlow) containing the stray frames received by that port. Stray frames are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Note: Stray Frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* results tab even if other flows in the same SmartFlow are listed.

Time	ILoad	FrameSize	PortName	FlowName	SourceIP	DestIP	RxFrames
12/4/01 7:22:21 PM	88.75	128	Port SMB 21 4	Flow A 1->2-1	192.85.1.3	192.85.2.4	44971
12/4/01 7:22:21 PM	88.75	128	Port SMB 21 4	Flow A 1->2-2	192.85.1.3	192.85.2.5	44971
12/4/01 7:22:21 PM	88.75	128	Port SMB 21 4	Flow A 1->2-3	192.85.1.3	192.85.2.6	44971
12/4/01 7:22:21 PM	88.75	128	Port SMB 21 4	Flow A 1->2-4	192.85.1.3	192.85.2.7	44971
12/4/01 7:22:34 PM	94.375	128	Port SMB 21 4	Flow A 1->2-0	192.85.1.3	192.85.2.3	47823
12/4/01 7:22:34 PM	94.375	128	Port SMB 21 4	Flow A 1->2-1	192.85.1.3	192.85.2.4	47823
12/4/01 7:22:34 PM	94.375	128	Port SMB 21 4	Flow A 1->2-2	192.85.1.3	192.85.2.5	47822
12/4/01 7:22:34 PM	94.375	128	Port SMB 21 4	Flow A 1->2-3	192.85.1.3	192.85.2.6	47822
12/4/01 7:22:34 PM	94.375	128	Port SMB 21 4	Flow A 1->2-4	192.85.1.3	192.85.2.7	47822
12/4/01 7:22:47 PM	97.1875	128	Port SMB 21 4	Flow A 1->2-0	192.85.1.3	192.85.2.3	49248

Figure 9-9. Throughput Stray Frames Results Tab

Each row on the *Stray Frames* results tab represents one flow of a SmartFlow. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

For a description of each field in Throughput stray frames results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Number of Records in Stray Frames Results

When *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. You can adjust this default limit by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

Chapter 10

Jumbo Test

This chapter provides information about the Jumbo test, which consists of the Frame Loss, Latency, and Latency Distribution tests.

In this chapter...

About the Jumbo Test

The Jumbo test is a combination of the Latency, Latency Distribution, and Frame Loss tests. The test measures latency variation (standard deviation) in addition to frame loss, latency, latency distribution, and sequencing. The Jumbo test updates all types of results in each test (except Latency SnapShot and Throughput) simultaneously.

For non-TeraMetrics receiving ports, the test calculates at different loads the minimum and maximum latency. (The exception is the LAN-3101A, which also calculates average latency). For TeraMetrics-based receiving ports, it calculates either the minimum and maximum latency, or the average and maximum latency. The test also provides information on the sequencing of frames under the varying conditions of latency exhibited in the test.

For information about each test, refer to the related chapter in this manual:

- *Chapter 11, "Frame Loss Test"*
- *Chapter 12, "Latency Test"*
- *Chapter 13, "Latency Distribution Test."*

Test Methodology

The Jumbo test checks packet tags to verify that packets transmitted by other SmartBits ports are received correctly, whether they are received in the correct sequence and whether there are missing packets. SmartFlow inspects the sequence tracking number in each frame to determine if frames were received in sequence.

For the duration of the test, the results show latency in microseconds for each group and/or flow received from each port transmitting to the receiving port. For non-TeraMetrics receiving ports, it shows the minimum, maximum, and average latency. For TeraMetrics-based receiving ports, it shows either the minimum and maximum latency or the average and maximum latency. It measures the distribution of up to eight latency bands. These results can be viewed by clicking the shortcut bar *Results* button, then clicking the *Latency Distribution* button.

Detailed Information on Result Calculations

For information about how latency is calculated, refer to *Chapter 12, "Latency Test."*

For a description of the two types of latency calculations (*cut-through* and *store and forward)* for the Latency test, see *"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349*.

For information about how standard deviation is calculated, see *"How Standard Deviation is Calculated" on page 361*.

For information on sequencing of frames and frame loss, see *"Sequence Tracking" on page 51*.

Interpreting Jumbo Test Results

Results of the Jumbo test appear in the main window at the conclusion of the test. You can view test results for this test in a chart, a summary table, or an itemized table. Results format can vary depending on whether or not you used frame size automation during the test. For more information, see *"Using Frame Size Automation (Global for All Flows)" on page 236*.

The standard deviation of latency distribution that appears in Jumbo test results is based on the latency distribution buckets in the Latency Distribution test. Standard deviation is also related to frame loss and starts increasing when frame loss begins to occur. If you look at Latency Distribution test results and see that the latency all falls into one bucket, the standard deviation is zero (0).

Figure 10-1 on page 332 shows an example of Jumbo test results in chart form, including the Latency, Latency Distribution, and Frame Loss tests results. Each of the other test results relates to the results that you see in the Jumbo test.

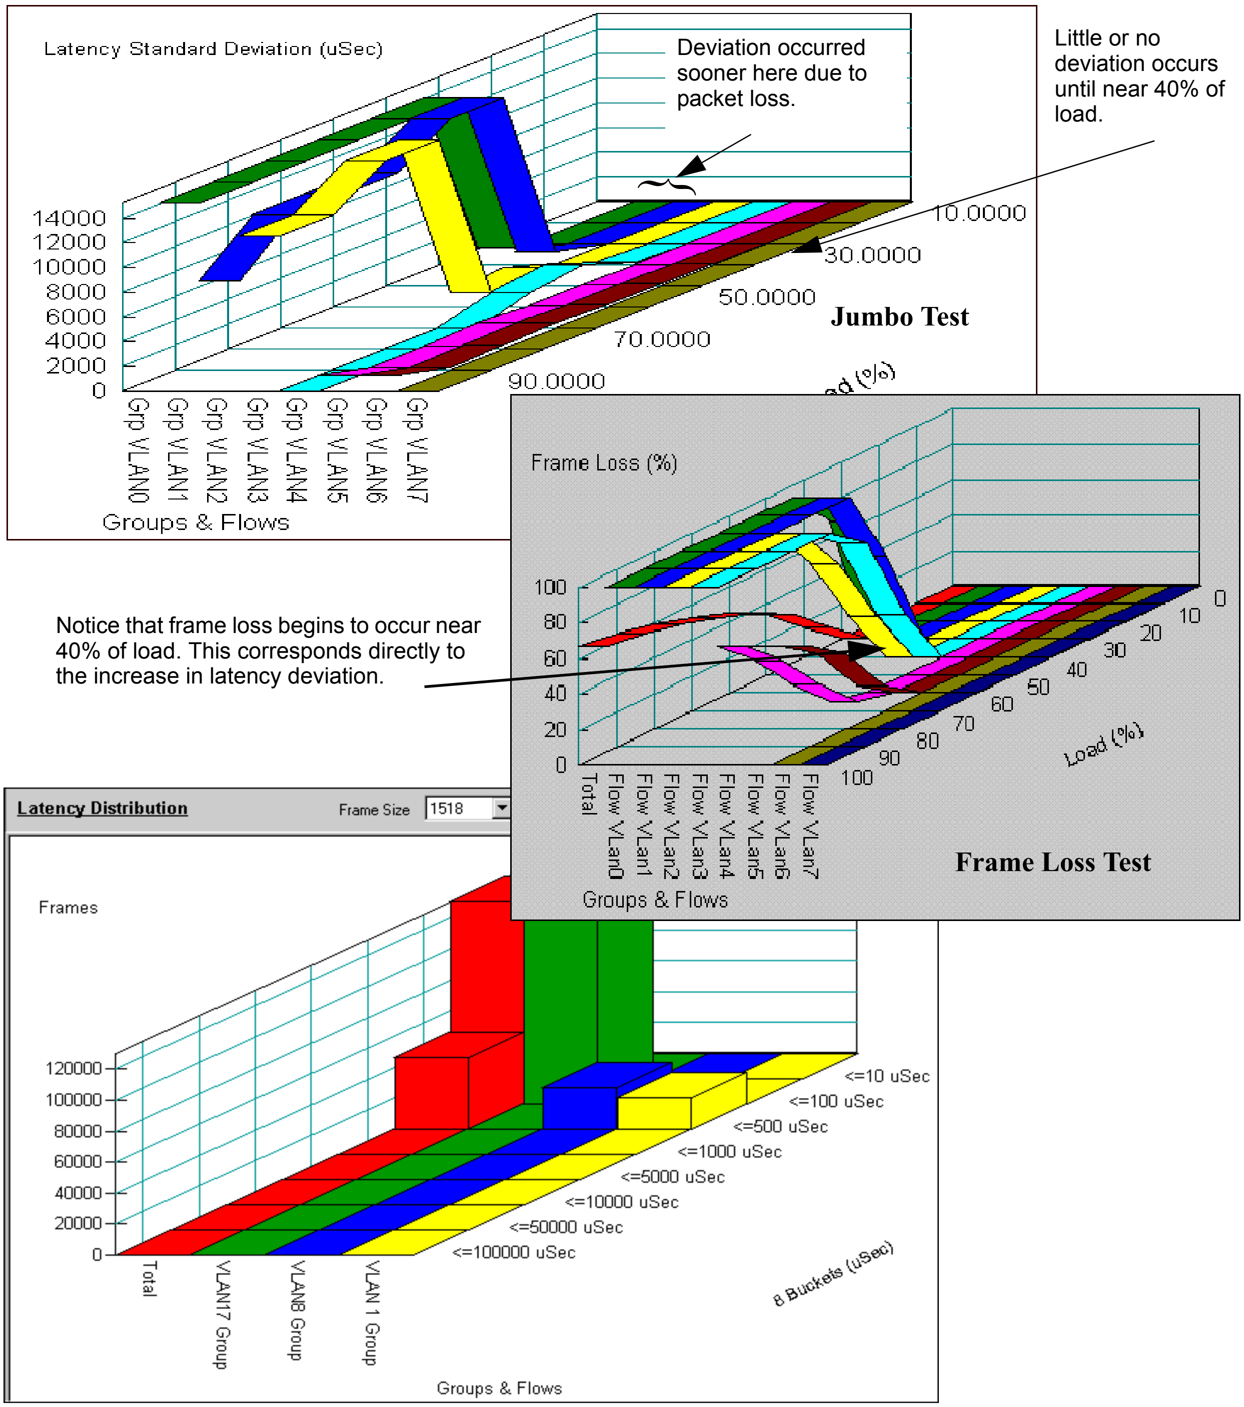

Figure 10-1. Jumbo Test Results

Chart (Jumbo)

The chart view (*Figure 10-2*) of the Jumbo test displays the number of frames delivered out of sequence for each flow and/or groups at a given load.

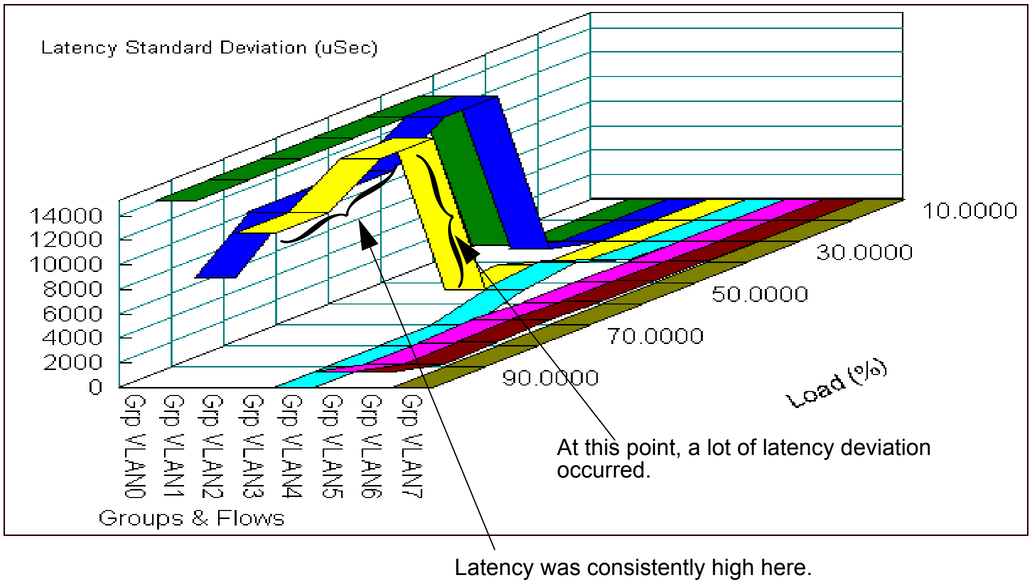

Figure 10-2. Jumbo Test Chart Results

Note: Right-click on the chart to access format options. The format options affect how the test information is displayed.

For a description of each field in Jumbo test chart results, refer to the SmartFlow online Help. Press **F1** over the tab or select **Help** from the menu bar.

Summary Results Table (Jumbo)

The Jumbo summary results table shows the number of frames delivered at a given load. *Figure 10-3* shows an example of the summary results in the Jumbo test.

Jumbo (in uSec) Frame Size 64

Name/Load	10.000%	20.000%	30.000%	40.000%	50.000%	60.000%	70.000%	80.000%	90.000%	100.000%
Total : Std. Dev.	120.70	120.70	120.80	321.20	327.80	336.60	342.10	345.40	347.00	347.40
Vlan1	122.00	122.00	122.00	121.90	54.50	7.60	6.00	5.40	5.20	4.90
Vlan 2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Vlan 3	115.80	115.60	115.20	7.00	5.30	4.80	4.50	4.40	4.30	4.30
Vlan 4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Figure 10-3. Jumbo Summary Tab

Note: Right-click to view frame size information.

For a description of each field in Jumbo test summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Detailed Table (Jumbo)

For flows and groups involved in a Jumbo test, a detailed table includes the latency and standard deviation at given loads. Totals for the load are displayed at the start of each new load percentage.

If you ran the test with custom frame sizes or loads, detailed test results display N/A for the frame size and load values, respectively.

You can view results according to the calculation for cut-through latency or store-and-forward latency. Use the drop-down list at the top of the page to select one type. See *"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349* for information on these options.

Notes:

- Your selection of latency option on the *Test Setup>Individual Tests* tab *(Min, Max, Sequencing* or *Average, Max*) determines what kinds of latency results are shown and whether sequence tracking information is available. For more information about these options, see *"Individual Tests Tab" on page 220*. For information about sequence tracking, *"Sequence Tracking" on page 51*.
- When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

Figure 10-4 on page 335 shows an example of detailed results for the Jumbo test. The test was run with both TeraMetrics and non-TeraMetrics cards and so shows both N/A and average latency values, respectively.

Name	Time	ILoad	TxFrames	RxFrames	LostFrames	StdDeviation	MinLatency	AveLatency	MaxLatency	InSequence
Total	11/16/01 2:28:54 PM	10.00000	1041600	1041600	0	120.70	11.30	58.92	324.20	1,041,600
Vlan1	11/16/01 2:28:54 PM	N/A	297600	297600	0	122.00	11.40	92.42	171.80	297,600
Vlan 2	11/16/01 2:28:54 PM	N/A	148800	148800	0	0.00	11.30	14.10	18.00	148,800
Vlan 3	11/16/01 2:28:54 PM	N/A	446400	446400	0	115.80	11.90	N/A	324.20	446,400
Vlan 4	11/16/01 2:28:54 PM	N/A	148800	148800	0	0.00	11.70	14.57	19.00	148,800
A 1-1->1-4	11/16/01 2:28:54 PM	10.00000	148800	148800	0	122.40	12.00	96.70	171.80	148,800
A 1-2->1-4	11/16/01 2:28:54 PM	10.00000	148800	148800	0	121.20	11.40	88.14	171.70	148,800
B 1-4->1-1	11/16/01 2:28:54 PM	5.00000	74400	74400	0	0.00	11.40	14.11	18.00	74,400
B 1-4->1-2	11/16/01 2:28:54 PM	5.00000	74400	74400	0	0.00	11.30	14.08	17.70	74,400
C 1-5->2-2	11/16/01 2:28:54 PM	10.00000	148800	148800	0	106.00	21.60	N/A	324.20	148,800
C 1-6->2-2	11/16/01 2:28:54 PM	10.00000	148800	148800	0	108.90	14.70	N/A	321.80	148,800
C 2-1->2-2	11/16/01 2:28:54 PM	10.00000	148800	148800	0	122.40	11.90	N/A	306.60	148,800
D 2-2->1-6	11/16/01 2:28:54 PM	5.00000	74400	74400	0	0.00	11.70	14.57	18.40	74,400
D 2-2->2-1	11/16/01 2:28:54 PM	5.00000	74400	74400	0	0.00	11.90	N/A	19.00	74,400
Total	11/16/01 2:29:21 PM	20.00000	2083200	2083200	0	120.70	11.30	59.11	338.00	2,083,200
Vlan1	11/16/01 2:29:21 PM	N/A	595200	595200	0	122.00	11.40	92.60	174.20	595,200
Vlan 2	11/16/01 2:29:21 PM	N/A	297600	297600	0	0.00	11.30	14.37	18.50	297,600
Vlan 3	11/16/01 2:29:21 PM	N/A	892800	892800	0	115.60	11.70	N/A	338.00	892,800
Vlan 4	11/16/01 2:29:21 PM	N/A	297600	297600	0	0.00	11.70	14.63	19.00	297,600
A 1-1->1-4	11/16/01 2:29:21 PM	20.00000	297600	297600	0	122.00	13.60	93.65	174.20	297,600
A 1-2->1-4	11/16/01 2:29:21 PM	20.00000	297600	297600	0	121.90	11.40	91.56	172.10	297,600
B 1-4->1-1	11/16/01 2:29:21 PM	10.00000	148800	148800	0	0.00	11.30	14.38	18.30	148,800
B 1-4->1-2	11/16/01 2:29:21 PM	10.00000	148800	148800	0	0.00	11.40	14.37	18.50	148,800

Figure 10-4. Jumbo Detail Tab with Both TeraMetrics and non-TeraMetrics Receiving Ports

Figure 10-5 shows an example of Jumbo detailed results for a test run with TeraMetrics modules. The *Min, Max, Sequencing* latency option was selected in the test setup (on the *Test Setup>Individual Tests* tab), so results display minimum and maximum latency. Average latency is displayed as **N/A**.

Name	FrameSiz	ILoad	TxFrames	RxFrames	LostFrames	StdDev	MinLatency	AveLatency	MaxLatency	InSeq	Out(
A 3A2->3A1-98	128	0.0333	140	140	0	0.00	154.60	N/A	263.30	140	
A 3A2->3A1-99	128	0.0333	140	140	0	0.00	154.40	N/A	268.50	140	
A 3A2->3A1-100	128	0.0333	140	140	0	0.00	154.60	N/A	265.20	140	
A 3A2->3A1-101	128	0.0333	140	140	0	0.00	154.50	N/A	263.30	140	
A 3A2->3A1-102	128	0.0333	140	140	0	0.00	154.70	N/A	263.50	140	
A 3A2->3A1-103	128	0.0333	140	140	0	0.00	154.50	N/A	268.50	140	
A 3A2->3A1-104	128	0.0333	140	140	0	0.00	154.60	N/A	263.30	140	
A 3A2->3A1-105	128	0.0333	140	140	0	0.00	154.50	N/A	265.50	140	
A 3A2->3A1-106	128	0.0333	140	140	0	0.00	154.50	N/A	263.50	140	
A 3A2->3A1-107	128	0.0333	140	140	0	0.00	154.60	N/A	268.20	140	
A 3A2->3A1-108	128	0.0333	140	140	0	0.00	154.60	N/A	263.50	140	

On TeraMetrics-based LAN and POS receiving ports, if you select *Min, Max, Sequencing* in the test setup, average latency is displayed as N/A.

Figure 10-5. Jumbo Detail Tab with TeraMetrics-based Receiving Ports

Figure 10-6 shows an example of Jumbo detail results for a test run with both TeraMetrics and non-TeraMetrics receiving ports and the *Min, Max, Sequencing* latency option selected on *Test Setup>Individual Tests* tab. Notice that for the TeraMetrics receiving ports, average latency is displayed as N/A whereas for the non-TeraMetrics LAN-3101A receiving ports, an average latency value is displayed.

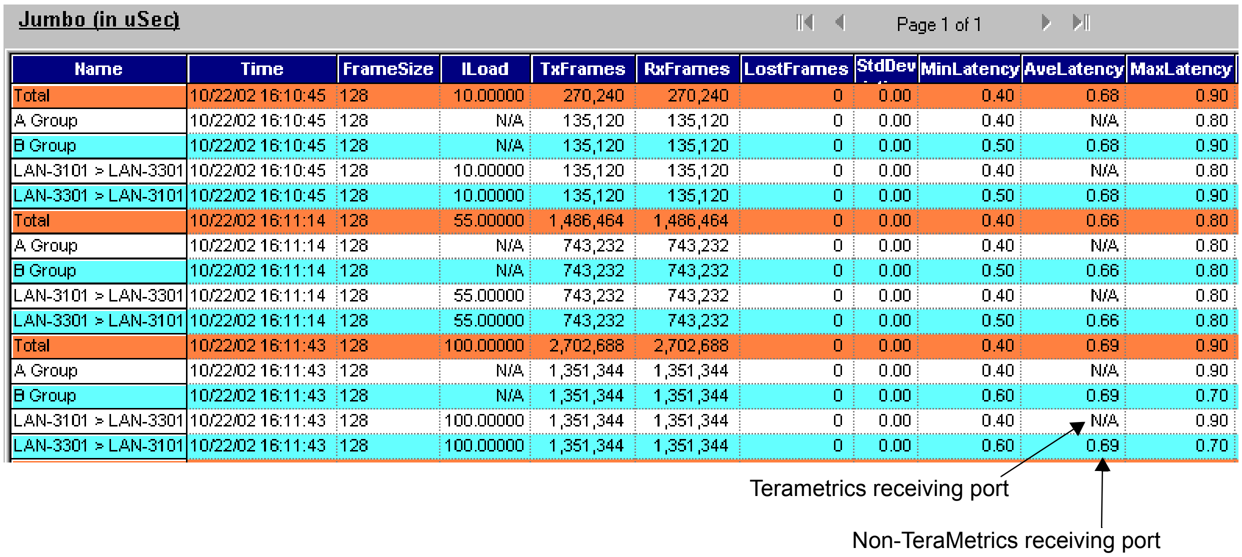

Jumbo (in uSec) — Page 1 of 1

Name	Time	FrameSize	ILoad	TxFrames	RxFrames	LostFrames	StdDev	MinLatency	AveLatency	MaxLatency
Total	10/22/02 16:10:45	128	10.00000	270,240	270,240	0	0.00	0.40	0.68	0.90
A Group	10/22/02 16:10:45	128	N/A	135,120	135,120	0	0.00	0.40	N/A	0.80
B Group	10/22/02 16:10:45	128	N/A	135,120	135,120	0	0.00	0.50	0.68	0.90
LAN-3101 > LAN-3301	10/22/02 16:10:45	128	10.00000	135,120	135,120	0	0.00	0.40	N/A	0.80
LAN-3301 > LAN-3101	10/22/02 16:10:45	128	10.00000	135,120	135,120	0	0.00	0.50	0.68	0.90
Total	10/22/02 16:11:14	128	55.00000	1,486,464	1,486,464	0	0.00	0.40	0.66	0.80
A Group	10/22/02 16:11:14	128	N/A	743,232	743,232	0	0.00	0.40	N/A	0.80
B Group	10/22/02 16:11:14	128	N/A	743,232	743,232	0	0.00	0.50	0.66	0.80
LAN-3101 > LAN-3301	10/22/02 16:11:14	128	55.00000	743,232	743,232	0	0.00	0.40	N/A	0.80
LAN-3301 > LAN-3101	10/22/02 16:11:14	128	55.00000	743,232	743,232	0	0.00	0.50	0.66	0.80
Total	10/22/02 16:11:43	128	100.00000	2,702,688	2,702,688	0	0.00	0.40	0.69	0.90
A Group	10/22/02 16:11:43	128	N/A	1,351,344	1,351,344	0	0.00	0.40	N/A	0.90
B Group	10/22/02 16:11:43	128	N/A	1,351,344	1,351,344	0	0.00	0.60	0.69	0.70
LAN-3101 > LAN-3301	10/22/02 16:11:43	128	100.00000	1,351,344	1,351,344	0	0.00	0.40	N/A	0.90
LAN-3301 > LAN-3101	10/22/02 16:11:43	128	100.00000	1,351,344	1,351,344	0	0.00	0.60	0.69	0.70

Figure 10-6. Latency in Jumbo Test Results for TeraMetrics and non-TeraMetrics Rx Ports

For a description of each field in the Jumbo test detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** over the tab or select **Help** from the main menu.

Stray Frames Results Table (Jumbo)

The Jumbo stray frames results table shows (per port) frames that were received but were not expected for the port. It also lists each flow (not SmartFlow) containing the stray frames received by that port. Stray frames are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Stray Frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* results tab, even if other flows in the same SmartFlow are listed.

Figure 10-7 shows an example of the *Stray Frames* results tab in the Jumbo test.

Jumbo (in uSec)

Time	ILoad	FrameS	PortName	FlowName	SourceIP	DestIP	RxFra	StdDev	MinLat	AveLat	MaxLa
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow NoVLAN 6A1-0->6A2-0	133.11.0.3	133.12.0.3	12476	0	0.4	N/A	0.8
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow NoVLAN 6A2-0->6A1-0	133.12.0.3	133.11.0.3	9510	0	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow VLAN 6A1-1-v111->6A2-1-v112	30.111.0.3	30.112.0.3	12475	0	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow VLAN 6A2-1-v112->6A1-1-v111	30.112.0.3	30.111.0.3	9509	0	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow NoVLAN 6A1-0->6A2-0	133.11.0.3	133.12.0.3	12476	0	0.4	N/A	0.8
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow NoVLAN 6A2-0->6A1-0	133.12.0.3	133.11.0.3	9510	0	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow VLAN 6A1-1-v111->6A2-1-v112	30.111.0.3	30.112.0.3	12475	0	0.4	N/A	0.8
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow VLAN 6A2-1-v112->6A1-1-v111	30.112.0.3	30.111.0.3	9509	0	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 6A1	Flow NoVLAN 1A2-0->1A1-0	133.2.0.3	133.1.0.3	9369	0.3	1.6	1.631	44.8
8/8/03 3:30:45 PM	10	128	Port SF 6A1	Flow NoVLAN 1A1-0->1A2-0	133.1.0.3	133.2.0.3	9369	0.7	1.7	1.731	60.6

Figure 10-7. Jumbo Stray Frames Tab

Each row in the *Stray Frames* results tab represents one flow of a SmartFlow. These items uniquely identify the flow: source IP address, destination IP address, and SmartFlow name.

Number of Records in Stray Frames Results

When *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. This default limit can be adjusted by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

For a description of each field in Jumbo test stray frames results, refer to the SmartFlow online Help. Press **F1** over the tab or select **Help** from the menu bar.

Bandwidth Chart Results (Jumbo)

The bandwidth chart view of the Jumbo test shows the bandwidth used or bit rate *per group* at each test load. It also shows the bit rates for the total transmit rate of all groups and the total receive rate of all groups at each test load.

Figure 10-8 shows an example of the *Bandwidth Chart* results tab in the Jumbo test.

Right-click anywhere on the tab to access the pop-up list. Select the unit (bit rate or frames) in which you want to view the results.

Note: The bandwidth chart supports up to 1,000 test load iterations. Any iterations greater than this amount are not displayed.

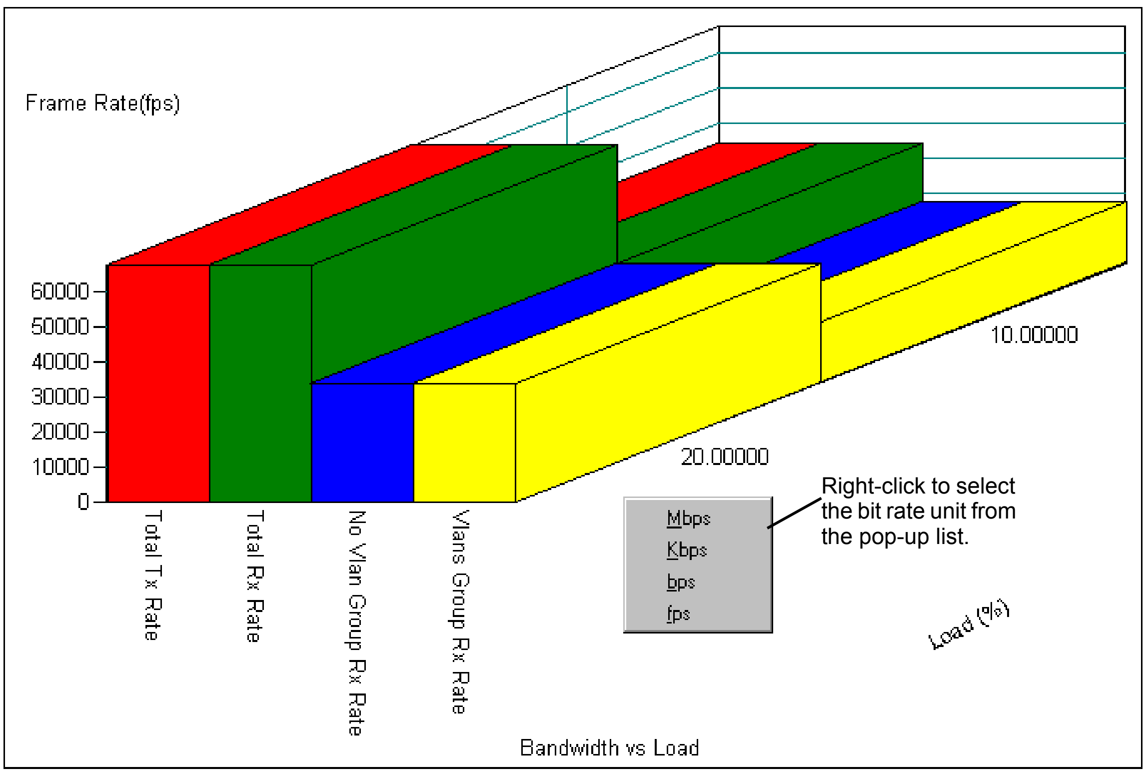

Figure 10-8. Bandwidth Chart - Jumbo Test

For a description of each field in Jumbo bandwidth chart results, refer to the SmartFlow online Help. Press **F1** over the tab or select **Help** from the menu bar.

Chapter 11

Frame Loss Test

In this chapter...

About the Frame Loss Test

The Frame Loss test measures the percentage of frames lost by the DUT that should have been forwarded. Sources of frame loss include:

- Network congestion
- Packets arriving late
- Out-of-sequence packets
- Frame errors.

Use this test to determine a DUT's ability to deliver frames in a sequenced flow of streams with specific routing priorities and at a stepped percentage of the wire rate.

Buffers on a DUT can dramatically minimize latency until the buffers are full. By using this test to determine the load at which the DUT drops frames, it can provide a valuable baseline for latency tests.

For information about how SmartFlow tracks packet sequence, see *"Sequence Tracking" on page 51*.

Test Methodology

The test runs according to the parameters specified for test duration, starting percentage load (based on the wire rate), a step percentage by which to increase the load during the test, and a stop percentage load at which the test ends. During the test, these values determine the overall duration for which the DUT is tested.

SmartFlow calculates frame loss as:

Frame Loss = Number of Frames Transmitted - Number of Frames Received

For each specified load percentage, the test progresses through these stages:

1. SmartFlow first sends a frame burst at the specified starting rate for the specified period of time.
2. After all frames are sent in all flows, SmartFlow queries the receiving ports to determine how many frames were received.
3. SmartFlow calculates total frame loss and out of sequence frames for each group and each flow in each group at that load percentage.
4. SmartFlow then steps up the load and sends a frame burst for the specified period of time. Steps 2 and 3 are repeated until the maximum load specified is reached.

For information on sequencing of frames and frame loss, see *"Sequence Tracking" on page 51*.

Interpreting Frame Loss Test Results

Results of the Frame Loss test automatically appear in the main window at the conclusion of the test. You can view test results for this test in graphical or tabular form. In tabular form, you can view a summary or detailed report.

Prioritization is not visible until there is congestion and thus frame loss. Notice in the test results that where frame loss starts occurring, the DUT also starts prioritizing how it handles groups and/or flows that it forwards.

Result formats can vary depending on whether or not you used frame size automation during the test. For more information, see *"Using Frame Size Automation (Global for All Flows)" on page 236*.

Chart Results (Frame Loss)

Frame Loss chart results illustrate a pattern of frame loss. If you used frame size automation, select the frame size at the top of the results tab for which you want to see results.

When a DUT treats multiple priorities the same, it can indicate the number of priority queues in the DUT. *Figure 11-1 on page 342* shows such an example of chart results for the Frame Loss test. In this example, eight flows were created, each with a different VLAN priority. Notice that the DUT is grouping flows with close priority together, such as VLAN7 and VLAN6, so that the eight priorities are being treated as four. This grouping indicates that the DUT has four priority queues.

Note: Right-click on the chart to access format options. The format options affect how the test information is displayed.

For a description of each field in Frame Loss test chart results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

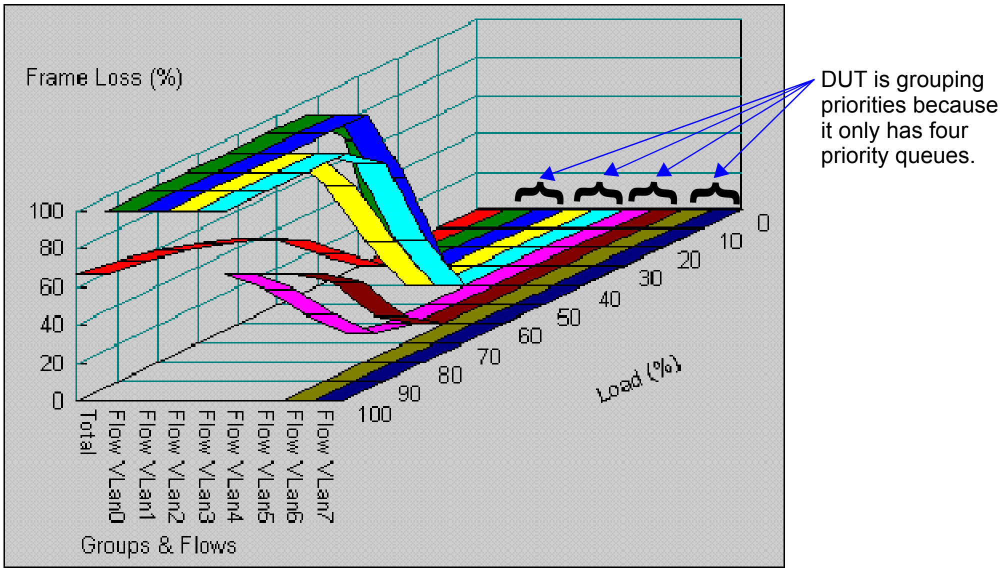

Figure 11-1. Frame Loss Chart Results

Summary Results Table (Frame Loss)

The Frame Loss summary results table shows the percentage of frames lost at the specified load intervals for flows and groups involved the test. If you used frame size automation, select the frame size at the top of the results tab for which you want to see results.

In the examples shown in *Figure 11-2* and *Figure 11-3 on page 343*, the DUT was configured to guarantee more bandwidth for high and medium traffic, but was also rate-limited in order not to exclude the low-priority traffic. The test included four ports. At 40%, frame loss starts to occur on two of the ports because the bandwidth was exceeded. These ports received 40% load from the three other ports, equaling 120% of load (which is greater than 100 Mbps).

Frame loss increases as traffic load increases.

No high priority traffic was dropped at any rate.

Name/Load	10.000%	20.000%	30.000%	40.000%	50.000%	60.000%	70.000%	80.000%	90.000%	100.000%
Total	0.00	0.00	0.00	7.23	14.45	23.98	30.78	35.88	39.85	42.85
Vlan1	0.00	0.00	0.00	0.00	0.29	16.96	28.86	37.79	44.74	50.00
Vlan 2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Vlan 3	0.00	0.00	0.00	16.86	33.53	44.64	52.58	58.53	63.16	66.66
Vlan 4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Figure 11-2. Frame Loss Summary Results

Notice that the high priority traffic (VLAN2 and VLAN4) had no traffic loss at any rate.

When frame loss is very small (less than 0.01%), the value 0.00 is shown in red. Click the cell to expand this value to show the exact percentage of frame loss.

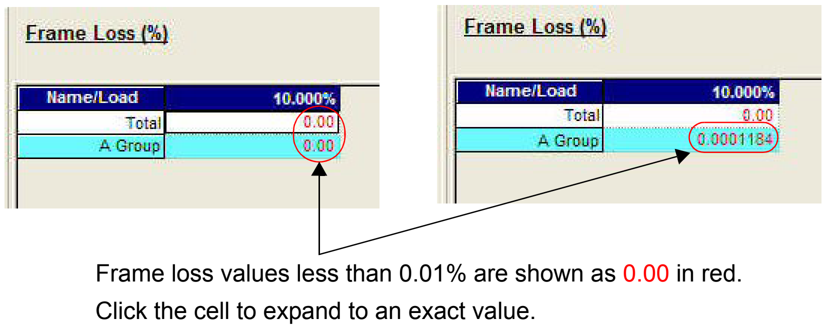

Figure 11-3. Frame Loss Summary Results

Note: Right-click to view frame size and frame loss information.

For a description of each field in Frame Loss test summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Detailed Results Table (Frame Loss)

For flows and groups involved in a frame loss test, the *Frame Loss detail* results tab shows the number of frames sent, received, lost, percent of loss, and out-of-sequence at a given load. Totals for the load are displayed at the start of each new load percentage. (See *Figure 11-4 on page 344*.)

Note: When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

When the number of frames sent is equal to the number received, there is no frame loss. As load increases, frame loss and frames received out of sequence also increase. If you ran the test with custom frame sizes or loads, detailed test results display N/A for the frame size and load values respectively.

If the test is run as part of the Jumbo test with the *Average, Max* latency option selected on the *Test Setup>Individual Tests* tab, results for the Jumbo Latency test show maximum and minimum latency instead of maximum and average latency. This is because minimum and maximum latency information is needed to track sequencing and report packet loss in the Frame Loss test. In addition, if you selected this option, sequencing statistics are not valid for any TeraMetrics and TeraMetrics-based POS and LAN cards. For details about these options, see *“Individual Tests Tab” on page 220*.

Totals for a given traffic load percentage and frame size

Number of frames sent minus received equals frames lost.

Name	Time	FrameSiz	Load	TxFrames	RxFrames	LostFram	Lost (%)
Total	11/16/01 2:34:27 PM	256	30.00000	951048	951048	0	0.00
Vlan1	11/16/01 2:34:27 PM	256	N/A	271728	271728	0	0.00
Vlan 2	11/16/01 2:34:27 PM	256	N/A	135864	135864	0	0.00
Vlan 3	11/16/01 2:34:27 PM	256	N/A	407592	407592	0	0.00
Vlan 4	11/16/01 2:34:27 PM	256	N/A	135864	135864	0	0.00
A 1-1->1-4	11/16/01 2:34:27 PM	256	30.00000	135864	135864	0	0.00
A 1-2->1-4	11/16/01 2:34:27 PM	256	30.00000	135864	135864	0	0.00
B 1-4->1-1	11/16/01 2:34:27 PM	256	15.00000	67932	67932	0	0.00
B 1-4->1-2	11/16/01 2:34:27 PM	256	15.00000	67932	67932	0	0.00
C 1-5->2-2	11/16/01 2:34:27 PM	256	30.00000	135864	135864	0	0.00
C 1-6->2-2	11/16/01 2:34:27 PM	256	30.00000	135864	135864	0	0.00
C 2-1->2-2	11/16/01 2:34:27 PM	256	30.00000	135864	135864	0	0.00
D 2-2->1-6	11/16/01 2:34:27 PM	256	15.00000	67932	67932	0	0.00
D 2-2->2-1	11/16/01 2:34:27 PM	256	15.00000	67932	67932	0	0.00
Total	11/16/01 2:34:54 PM	256	40.00000	1268064	1177252	90,812	7.16
Vlan1	11/16/01 2:34:54 PM	256	N/A	362304	362304	0	0.00
Vlan 2	11/16/01 2:34:54 PM	256	N/A	181152	181152	0	0.00
Vlan 3	11/16/01 2:34:54 PM	256	N/A	543456	452644	90,812	16.71
Vlan 4	11/16/01 2:34:54 PM	256	N/A	181152	181152	0	0.00
A 1-1->1-4	11/16/01 2:34:54 PM	256	40.00000	181152	181152	0	0.00
A 1-2->1-4	11/16/01 2:34:54 PM	256	40.00000	181152	181152	0	0.00
B 1-4->1-1	11/16/01 2:34:54 PM	256	20.00000	90576	90576	0	0.00
B 1-4->1-2	11/16/01 2:34:54 PM	256	20.00000	90576	90576	0	0.00
C 1-5->2-2	11/16/01 2:34:54 PM	256	40.00000	181152	137832	43,320	23.91
C 1-6->2-2	11/16/01 2:34:54 PM	256	40.00000	181152	157753	23,399	12.91
C 2-1->2-2	11/16/01 2:34:54 PM	256	40.00000	181152	157059	24,093	13.29

Figure 11-4. Frame Loss Detailed Results

As load increases, frame loss and frames received out of sequence also increase.

The number of out-of-sequence frames can be due to frame loss or to frame reordering by the device. Any frame loss causes frames to be received out of sequence, so no conclusions can be drawn from the sequence results if there is frame loss.

Examples:
Frames are sent out in the order 1, 2, 3, 4, 5.
Frames are received in the order of 1, 2, 5, 3, 4. This yields two out-of-sequence frames.

Frames are sent out in the order 1, 2, 3, 4, 5.
Frames are received in the order of 1, 2, 4, 3, 5. This yields three out-of-sequence frames.

Frames are sent out in the order 1, 2, 3, 4, 5, 6, 7.
Two frames are lost, so frames are received in the order 1, 2, 4, 6, 7. This yields two out-of-sequence frames.

See *"Sequence Tracking" on page 51* for more information.

For a description of each field in the Frame Loss test detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Stray Frames Results Table (Frame Loss)

The Frame Loss stray frames results table shows *per port* the frames that were received but were not destined (expected) for that port. It also lists each flow (not SmartFlow) containing the stray frames received by that port. Stray frames are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Stray frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* results tab even if other flows in the same SmartFlow are listed.

Figure 11-5 shows an example of stray frames results for the Frame Loss test.

Time	ILoad	FrameSize	PortName	FlowName	SourceIP	DestIP	RxFrames
12/4/01 7:20:46 PM	50	0	Port SMB 21 4	Flow A 1->2-0	192.85.1.3	192.85.2.3	25335
12/4/01 7:20:46 PM	50	0	Port SMB 21 4	Flow	0.0.0.0	0.0.0.0	0
12/4/01 7:20:46 PM	50	0	Port SMB 21 4	Flow SMB 21 1	0.0.0.0	0.0.0.0	0
12/4/01 7:20:46 PM	50	0	Port SMB 21 4	Flow	0.0.0.0	0.0.0.0	0
12/4/01 7:20:46 PM	50	0	Port SMB 21 4	Flow	0.0.0.0	0.0.0.0	0
12/4/01 7:21:00 PM	100	0	Port SMB 21 4	Flow A 1->2-0	192.85.1.3	192.85.2.3	50674
12/4/01 7:21:00 PM	100	0	Port SMB 21 4	Flow	0.0.0.0	0.0.0.0	0
12/4/01 7:21:00 PM	100	0	Port SMB 21 4	Flow SMB 21 1	0.0.0.0	0.0.0.0	0
12/4/01 7:21:00 PM	100	0	Port SMB 21 4	Flow	0.0.0.0	0.0.0.0	0
12/4/01 7:21:00 PM	100	0	Port SMB 21 4	Flov	0.0.0.0	0.0.0.0	0

Figure 11-5. Frame Loss Stray Frames Tab

Each row in the *Stray Frames* results tab represents one flow of a SmartFlow. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

Number of Records in Stray Frames Results

When the *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. You can adjust this default limit by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

For a description of each field in Frame Loss test stray frames results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Chapter 12

Latency Test

In this chapter...

About the Latency Test

After you have determined load tolerance limits by using the Frame Loss test, run the Latency test to measure the latency (transmission delay) introduced by the DUT. It is recommended that you run the Latency test at the maximum rate at which no packet loss occurs (the throughput rate per RFC 2544) or at a lower rate.

The Latency test helps you determine:

- How a DUT delivers traffic across multiple destination ports.
- How the latency varies.
- Latency per port.
- A device's latency characteristics.
- The effect on latency when you modify the DUT's RAM and re-run the test.

The Latency test measures latency for received frames only. Results that display include latency for each trial; frame length per port; and the minimum, maximum, and average latency per port.

For SmartFlow to determine latency correctly, the DUT's egress port should be able to forward packets without buffering (queuing). Otherwise, the latency measurements include buffering delays.

Tip: Because the DUT may be buffering packets when the maximum rate with no packet loss is reported, it may be more accurate to run the Latency test at a slightly lower rate than the throughput rate.

For non-TeraMetrics receiving ports, the test calculates the minimum and maximum latency at different loads. (The exception is the LAN-3101A, which also calculates average latency.) For TeraMetrics-based receiving ports, it calculates either the minimum and maximum latency, or the average and maximum latency. See *Table 12-2, "How Latency is Displayed, Depending on Flow's Receiving Port and Test Options," on page 353* for more information.

In SmartFlow, latency is defined as the time it takes the DUT to forward a packet from one SmartBits port to another SmartBits port. In general, factors that can affect latency include:

- Hop counts
- Number of ISPs traversed
- Network congestion
- The DUT's internal clock tolerance limits, plus variations between the DUT's clock and the chassis
- A busy router or switch taking longer to look up the routing (address) table
- Transit or processing buffers in the DUT.

Test Methodology

This test tracks the latency of each test frame on a frame-by-frame basis. SmartFlow reads the time that the sending SmartBits port sent the frame (transmit timestamp) and the time that the receiving SmartBits recognizes the trigger frame, which is the receive timestamp.

Latency is calculated as:

Latency = (Receive Timestamp) – (Transmit Timestamp)

You can set up multiple unicast flows for each port, some unidirectional and some bi-directional, each with a different frame size, protocol, and scheme for priority routing.

The test uses the specified test duration, a starting percentage load (based on the wire rate), a step percentage by which to increase the load during the test, and a stop percentage load at which the test ends. During the test, these values determine the duration for which the DUT is tested at a specific load.

Note: Latency values are only for the frames that were not dropped. If all frames are dropped for a receiving port, no results are displayed for that port. If all frames are dropped for all receiving ports, the results tabs are blank.

Latency Measurements for Cut-through and Store-and-Forward Devices

When testing with TeraMetrics-based LAN modules, you can view Latency test results based on either of these two alternative calculations:

- *Cut-through Latency:* This calculation is appropriate when testing *bit-forwarding* devices or functionality. In this calculation, latency values are based on the last bit out the transmitting SmartBits port and last bit in the receiving SmartBits port.
- *Store-and-Forward Latency:* This calculation is appropriate when testing *store-and-forward* devices or functionality. In this calculation, latency values are based on the last bit out the transmitting SmartBits port and first bit in the receiving SmartBits port after the preamble (LIFO).

For both the Latency test and Jumbo test (which includes the Latency test), SmartFlow calculates both types of results automatically. You can select the desired type from a pull-down list box at the top of the *Detail* results tab (for the Latency test or Jumbo test). (See *Figure 12-1 on page 350*.) When you select a type on one tab, the results on related tabs are updated to the selected type.

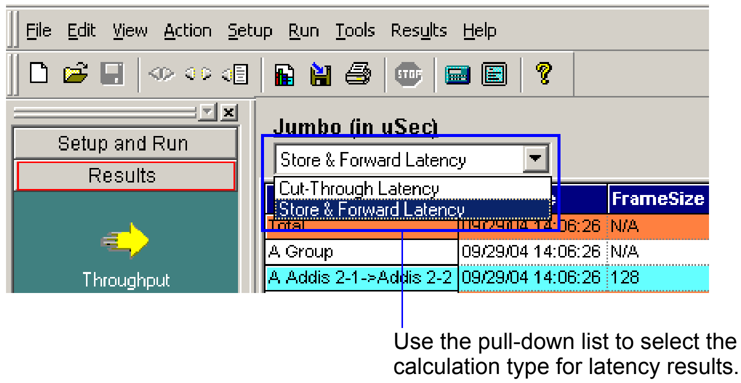

Figure 12-1. Selecting Cut-through or Store & Forward Latency Values

Note: When your test uses different port types, the transmit and receive speed for a flow can differ. In this case, the bit time for the transmitting SmartBits port is used in calculating the store-and-forward latency.

For additional information about latency calculations, refer to the SmartFlow online Help. Press **F1** over the Results page or select **Help** from the main menu.

How SmartFlow Calculates Store-and-Forward Latency

By default, SmartFlow calculates latency according to the cut-through (bit-forwarding) method. This is one of two possible ways to measure latency, according to RFC 1242. The RFC states:

- For bit-forwarding devices, latency is “the time interval starting when the end of the first bit of the input frame reaches the input port and ending when the start of the first bit of the output frame is seen on the output port.”
- For store and forward devices, latency is “the time interval starting when the last bit of the input frame reaches the input port and ending when the first bit of the output frame is seen on the output port.”

To determine store-and-forward latency for received packets, SmartFlow subtracts the bit time, derived from the frame length and transmit speed, from the cut-through latency measurement.

Example: Assume the packet length is 64 bytes and the transmit speed is 10 Mbps.

The number of bits transferred is 64 X 8 = 512 bits. The 8-byte preamble is not included in this calculation. At 10 Mbps, the bit rate is 100 nanoseconds. At 100 Mbps, it is 10 nanoseconds. This gives:

512 bits X 100 nsec = 51200 nsec for 10 MB port speeds (51.2 us)
512 bits X 10 nsec = 5120 nsec for 100 MB port speeds (5.12 us).

To obtain store-and-forward latency result values, SmartFlow subtracts 5.12 us from the cut-through result values.

Table 12-1 shows how this calculation applies to different frame sizes, based on a port speed of 100 Mbps. It lists the time, in microseconds (us), that SmartFlow subtracts from its cut-through latency result values to determine the store-and-forward latency result values.

Table 12-1. Subtraction Times from Cut-through for Store-and-Forward Latency

If frame size is ...	SmartFlow subtracts ... (us)[1]
64	5.12
128	10.24
256	20.48
512	40.96
1280	102.40
1518	121.44
9018	721.44

1 Values are based on 100 Mbps port speed.

Interpreting Latency Test Results

Results of the Latency test automatically appear in the main window at the conclusion of the test. You can view test results for this test in graphical or tabular form. In tabular form, you can view a summary or detailed report.

If you are tracking the DUT's priority handling, look for prioritization occurring where large latency occurs. Priority takes effect on latency even before the point of congestion. Sharp latency increases are due to prioritization occurring from increasing traffic load.

Note: Results formats can vary depending on whether or not frame size automation was used during the test. For more information, see *"Using Frame Size Automation (Global for All Flows)" on page 236*.

Various factors affect how latency is calculated and displayed in results:

- Type of receiving port for the flow (TeraMetrics or non-TeraMetrics-based).
- Whether the test was run alone or as part of the Jumbo test.
- Which latency option (for TeraMetrics-based ports) was selected in test setup.
- Which latency calculation you selected: cut-through or store-and-forward (*"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349*).
- Whether groups contain flows with "mixed" receiving ports (both TeraMetrics-based and non-TeraMetrics-based).

Table 12-2 on page 353 shows the various ways latency may be calculated and displayed in results.

Table 12-2. How Latency is Displayed, Depending on Flow's Receiving Port and Test Options

If flow's receiving ports are ...	and Latency test is run ...	Latency for the flow is displayed as ...
Non-TeraMetrics-based	Alone or part of Jumbo test	**Chart/Summary:** Average or maximum **Detail:** Minimum, average, and maximum
TeraMetrics and TeraMetrics-based only	Alone	**Chart/Summary:** Average or maximum **Detail:** Average and maximum or minimum and maximum, depending on test setup option
TeraMetrics and TeraMetrics-based only	Part of Jumbo test	**Chart/Summary:** Maximum **Detail:** Average and maximum or minimum and maximum, depending on test setup option
Both non-TeraMetrics, and TeraMetrics or TeraMetrics-based (*separate group for each port type*)	Part of Jumbo test	*For non-TeraMetrics Rx ports:* **Chart/Summary:** Average or maximum **Detail:** Minimum, average, and maximum *For TeraMetrics or TeraMetrics-based Rx ports:* **Chart/Summary:** If maximum is selected, maximum. If average is selected, then no display in chart and N/A in summary. **Detail:** Average and maximum or minimum and maximum, depending on test setup option
Both non-TeraMetrics and TeraMetrics or TeraMetrics-based (*mixed ports within same groups*)	Part of Jumbo test	*For non-TeraMetrics Rx ports:* **Chart/Summary:** Average or maximum If average is selected, average is *only* for non-TeraMetrics ports. **Detail:** Minimum, average, and maximum *For TeraMetrics or TeraMetrics-based Rx ports:* **Chart/Summary:** If maximum is selected, maximum. If average is selected, average is *only* for *non*-TeraMetrics ports. **Detail:** Average and maximum or minimum and maximum, depending on test setup option

Chart Results (Latency)

The Latency chart results graph the pattern of latency for every group and/or flow in the test at each traffic load percentage. If you used frame size automation, select the frame size at the top of the results tab for which you want to see results.

You can view chart results by average latency or maximum latency. Right-click on the graph to display the viewing options menu.

You can view results according to the calculation for cut-through latency or store-and-forward latency. Use the drop-down list at the top of the page to select one type. See *"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349* for information on these options.

If the Latency test was run as part of the Jumbo test, with all TeraMetrics and TeraMetrics-based receiving ports and with the *Maximum latency* right-click option selected, average latency is not displayed at all in results. (See *Figure 12-2*.)

Note: Rght-click on the chart to access format options. The format options affect how the test information is displayed.

For field descriptions for Summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

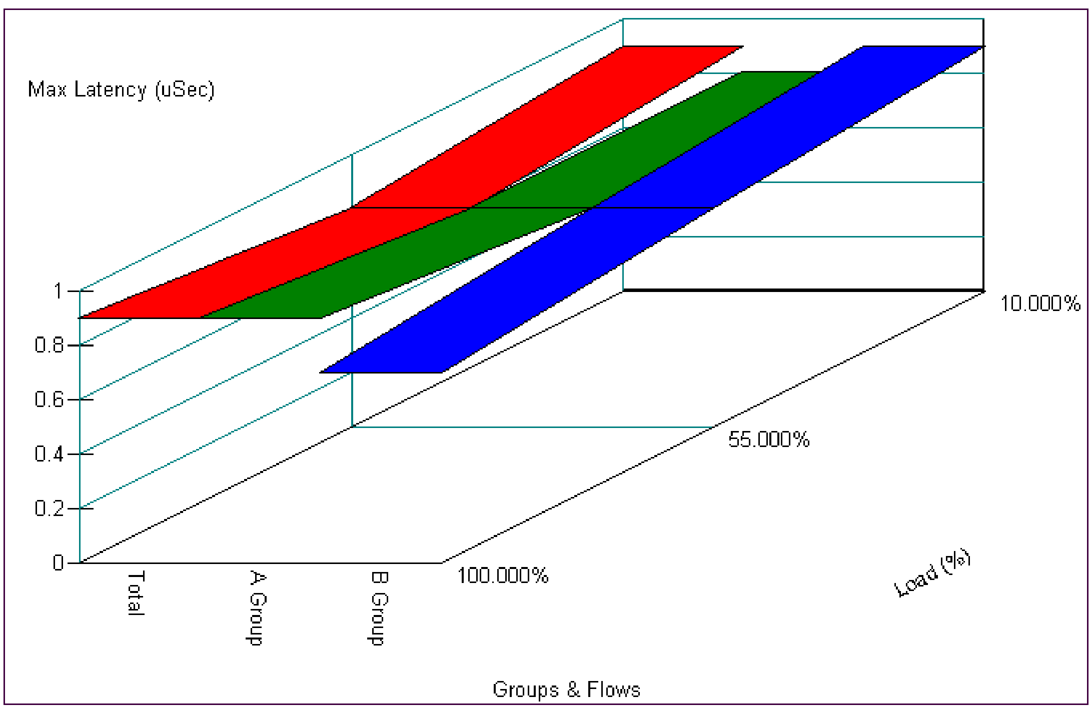

Figure 12-2. Latency Test Chart Results Shows Max Latency When Run as Part of Jumbo Test

In *Figure 12-3 on page 355*, the right-click option selected is *Average latency.* All the receiving ports in Group A are TeraMetrics ports, so that group is not displayed.

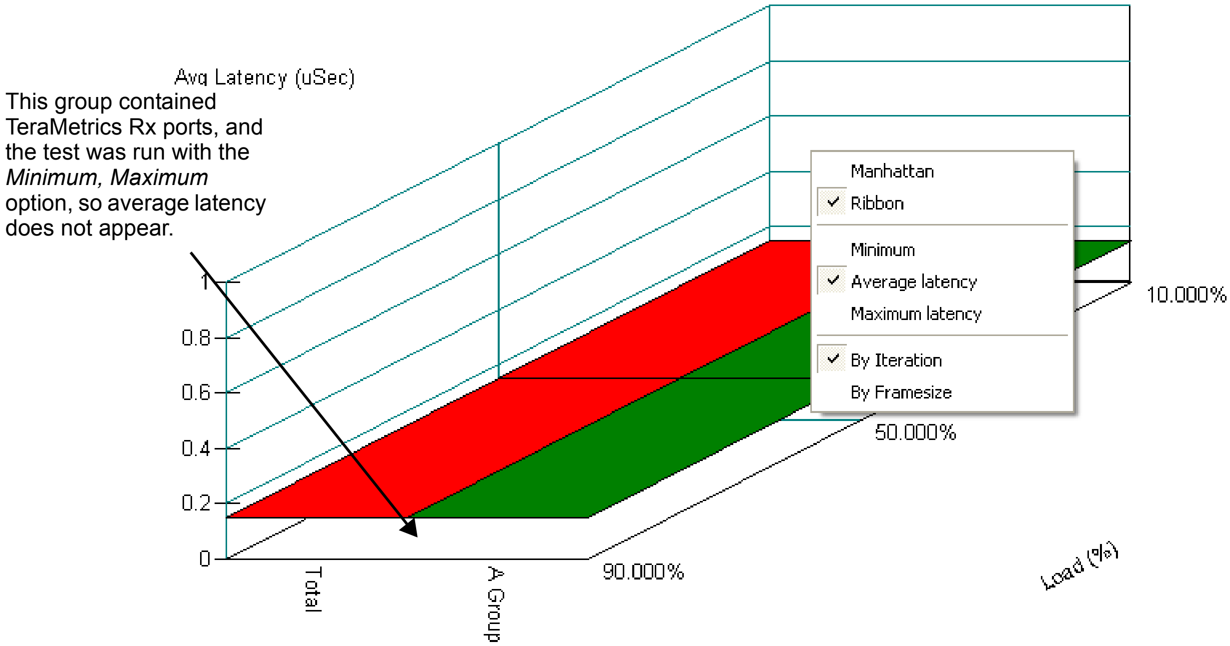

Figure 12-3. Latency Chart Results for TeraMetrics and Non-TeraMetrics Ports and Min/Max Latency Option

Summary Results Table (Latency)

The Latency summary results table shows the latency for each group and/or flow at every traffic load. As shown in the example in *Figure 12-4*, latency increases as the load increases.

If frame size automation was used, select the frame size for which you want to see results at the top of the results tab.

You can view summary results by either average latency or maximum latency. Right-click on the summary tab to display the viewing options menu. If the Latency test was run as part of the Jumbo test, only maximum latency is available as a display option.

You can view results according to the calculation for cut-through latency or store-and-forward latency. Use the drop-down list at the top of the page to select one type. See *"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349* for information on these options.

Latency usually increases as traffic load increases.

Note: Right-click to view frame size and minimum/average/maximum latency information. Click the required option. (Minimum/average/maximum values are the minimum/average/maximum of all minimums/averages/maximums for frames of the noted size.)

For field descriptions for Summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Name/Load	10.000%	20.000%	30.000%	40.000%	50.000%	60.000%	70.000%	80.000%	90.000%	100.000%
Total	58.92	59.11	58.70	58.97	430.29	418.73	399.91	379.54	358.45	339.54
Vlan1	92.42	92.60	91.94	92.46	742.67	783.96	806.01	819.70	825.33	827.18
Vlan 2	14.10	14.37	14.19	14.13	15.05	14.15	14.66	14.30	14.21	14.18
Vlan 3	324.20	338.00	338.50	831.40	832.30	831.90	832.40	832.10	836.80	834.90
Vlan 4	14.57	14.63	14.72	14.72	14.91	14.72	14.88	14.80	14.93	14.90

Figure 12-4. Latency Summary Results (Cut-through Calculation)

Detailed Table (Latency)

Latency detail results provide statistics for group and/or flow at a given load.

- For non-TeraMetrics-based receiving ports, they show the minimum, average, and maximum latency.
- For TeraMetrics-based receiving ports, they show either minimum and maximum or average and maximum latency.

You can view results according to the calculation for cut-through latency or store-and-forward latency. Use the drop-down list at the top of the page to select one type. See *"Latency Measurements for Cut-through and Store-and-Forward Devices" on page 349* for information on these options.

Note: When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

Totals for the load are displayed at the start of each new load percentage. (See *Figure 12-5*).

Name	Time	FrameSi	ILoad	TxFram	RxFram	MinLat	AveLat	MaxLat	Tx fps
Total	08/08/03 14:35:39	N/A	10.000	101,344	101,34	14.20	15.808	18.80	33,781
No Vlan Group	08/08/03 14:35:39	N/A	N/A	50,672	50,672	14.80	16.155	18.80	16,891
Vlans Group	08/08/03 14:35:39	N/A	N/A	50,672	50,672	14.20	15.461	17.80	16,891
No Vlan 1A1-0->1	08/08/03 14:35:39	128	5.0000	12,668	12,668	16.20	N/A	18.80	4,223
No Vlan 1A2-0->1	08/08/03 14:35:39	128	5.0000	12,668	12,668	15.60	N/A	16.60	4,223
No Vlan 6A1-0->6	08/08/03 14:35:39	128	5.0000	12,668	12,668	15.10	16.264	17.30	4,223
No Vlan 6A2-0->6	08/08/03 14:35:39	128	5.0000	12,668	12,668	14.80	16.046	17.30	4,223
Vlans 1A1-1-v101-	08/08/03 14:35:39	128	5.0000	12,668	12,668	15.10	N/A	17.80	4,223
Vlans 1A2-1-v102-	08/08/03 14:35:39	128	5.0000	12,668	12,668	15.00	N/A	16.60	4,223
Vlans 6A1-1-v111-	08/08/03 14:35:39	128	5.0000	12,668	12,668	14.20	15.413	16.60	4,223
Vlans 6A2-1-v112-	08/08/03 14:35:39	128	5.0000	12,668	12,668	14.30	15.508	16.70	4,223
Total	08/08/03 14:36:04	N/A	20.000	202,696	202,69	14.30	15.763	18.90	67,565
No Vlan Group	08/08/03 14:36:04	N/A	N/A	101,348	101,34	14.80	16.058	18.90	33,783
Vlans Group	08/08/03 14:36:04	N/A	N/A	101,348	101,34	14.30	15.468	18.00	33,783
No Vlan 1A1-0->1	08/08/03 14:36:04	128	10.000	25,337	25,337	15.90	N/A	18.90	8,446
No Vlan 1A2-0->1	08/08/03 14:36:04	128	10.000	25,337	25,337	15.60	N/A	18.00	8,446
No Vlan 6A1-0->6	08/08/03 14:36:04	128	10.000	25,337	25,337	14.80	16.063	17.30	8,446
No Vlan 6A2-0->6	08/08/03 14:36:04	128	10.000	25,337	25,337	14.80	16.053	17.30	8,446
Vlans 1A1-1-v101-	08/08/03 14:36:04	128	10.000	25,337	25,337	15.00	N/A	18.00	8,446
Vlans 1A2-1-v102-	08/08/03 14:36:04	128	10.000	25,337	25,337	15.00	N/A	16.70	8,446
Vlans 6A1-1-v111-	08/08/03 14:36:04	128	10.000	25,337	25,337	14.50	15.420	16.70	8,446
Vlans 6A2-1-v112-	08/08/03 14:36:04	128	10.000	25,337	25,337	14.30	15.515	16.70	8,446

Total number of frames received for all flows at this traffic load

Non-TeraMetrics-based ports

With the *Minimum, Maximum* option selected, TeraMetrics ports display N/A.

Figure 12-5. Latency Detail Results (Cut-Through Latency)

Scroll to the right to see all of the results columns on your monitor.

For field descriptions of the Detail results as well as the customized detailed reporting function and cut-through versus store-and-forward latency calculations, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Stray Frames Results Table (Latency)

The Latency stray frames results table shows *per port* the frames that were received but were not destined (expected) for that port. It also lists each flow (not SmartFlow) containing the stray frames received by that port. Stray frames (*Figure 12-6*) are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Stray frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* results tab even if other flows in the same SmartFlow are listed.

Latency (in uSec)

Time	ILoad	FrameS	PortName	FlowName	SourceIP	DestIP	RxFrames	MinLat	AveLat	MaxLat
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow NoVLAN 6A1-0->6A2-0	133.11.0.3	133.12.0.3	12476	0.4	N/A	0.8
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow NoVLAN 6A2-0->6A1-0	133.12.0.3	133.11.0.3	9510	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow VLAN 6A1-1-v111->6A2-1-v112	30.111.0.3	30.112.0.3	12475	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow VLAN 6A2-1-v112->6A1-1-v111	30.112.0.3	30.111.0.3	9509	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow NoVLAN 6A1-0->6A2-0	133.11.0.3	133.12.0.3	12476	0.4	N/A	0.8
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow NoVLAN 6A2-0->6A1-0	133.12.0.3	133.11.0.3	9510	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow VLAN 6A1-1-v111->6A2-1-v112	30.111.0.3	30.112.0.3	12475	0.4	N/A	0.8
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow VLAN 6A2-1-v112->6A1-1-v111	30.112.0.3	30.111.0.3	9509	0.4	N/A	0.9
8/8/03 3:30:45 PM	10	128	Port SF 6A1	Flow NoVLAN 1A2-0->1A1-0	133.2.0.3	133.1.0.3	9369	1.6	1.631	44.8
8/8/03 3:30:45 PM	10	128	Port SF 6A1	Flow NoVLAN 1A1-0->1A2-0	133.1.0.3	133.2.0.3	9369	1.7	1.731	60.6

Figure 12-6. Latency Stray Frames Tab

Each row in the *Stray Frames* results tab represents one flow of a SmartFlow. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

Number of Records in Stray Frames Results

When the *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. You can adjust this default limit by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

For a description of each field in Latency test stray frames results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Chapter 13

Latency Distribution Test

In this chapter...

About the Latency Distribution Test

Use the Latency Distribution test to examine latency across a user-defined distribution of eight possible latency time buckets. For each latency bucket defined on the *Test Setup>Individual Tests* tab, the test shows the number of packets received whose latency falls within the range of that bucket.

This test can be run at whatever loads need to be tested. Compared to the Latency test, this test can provide a more detailed view of latency behavior at the DUT load tolerance limits.

Test Methodology

This test tracks the latency of each test frame on a frame-by-frame basis and places latency results into eight latency bands (buckets). To determine latency, SmartFlow reads the time that the sending SmartBits port sent the frame (transmit timestamp) and the time that the receiving SmartBits port recognizes the receive timestamp in the frame.

To select bucket sizes that are most meaningful, first run the Latency test to find the minimum and maximum latency. Then distribute the buckets evenly over that time.

How Latency is Calculated

Latency is calculated as:

Latency = (Receive Timestamp) - (Transmit Timestamp)

Multiple unicast flows can be set up for each port, some unidirectional and some bi-directional, each with a different frame size, protocol, and scheme for priority routing.

The test uses the specified test duration, a starting percentage load (based on the wire rate), a step percentage by which to increase the load during the test, and a stop percentage load at which the test ends. During the test, these values determine the duration for which the DUT is tested at a specific load.

Note: Latency values are only for the frames that were not dropped.

What is Average Latency and Standard Deviation?

Measurements of *central tendency*, such as average latency and standard deviation, are measures of the location of the middle or the center of a distribution. If the number of received packets is plotted against the latency for each packet, you typically get a normal distribution. The mean of latency distribution of all received packets is near the center of this distribution. This does not mean that all of the packets have the same latency. The latency of each individual packet varies somewhat from the mean.

Standard deviation is a statistical way of calculating how much the latency of a set of received packets varies from the mean. When individual latencies of received packets vary greatly from the mean, the latency has a high standard deviation. When latency does not vary greatly from the mean, the latency has a low standard deviation.

How Standard Deviation is Calculated

Standard deviation is calculated using the steps described in this section.

These are the symbols used in the standard deviation calculation:

i – any integer, such as 1, 2, 3

X – In probability theory, this represents a discrete random variable.
In SmartFlow, it represents latency.

x_i – A discrete value that X can take on.
In SmartFlow, it is the latency value that represents bucket i. Here is how x_i is calculated:

$$x_i = \frac{b_i - b_{i-1}}{2}$$

where b_i is the user-defined upper boundary of the latency bucket specified on the *Test Setup>Individual Tests* tab. The exception is b_0 (which is zero microsecond)s and the lower boundary of the first bucket. See *Figure 13-1*.

f_i – In probability theory, the occurrence of the X that takes on x_i.
In SmartFlow, it is the number of packets that fall into the latency bucket i.

n –In probability theory, the number of discrete values that X can take on.
In SmartFlow, it is the number of latency buckets (8).

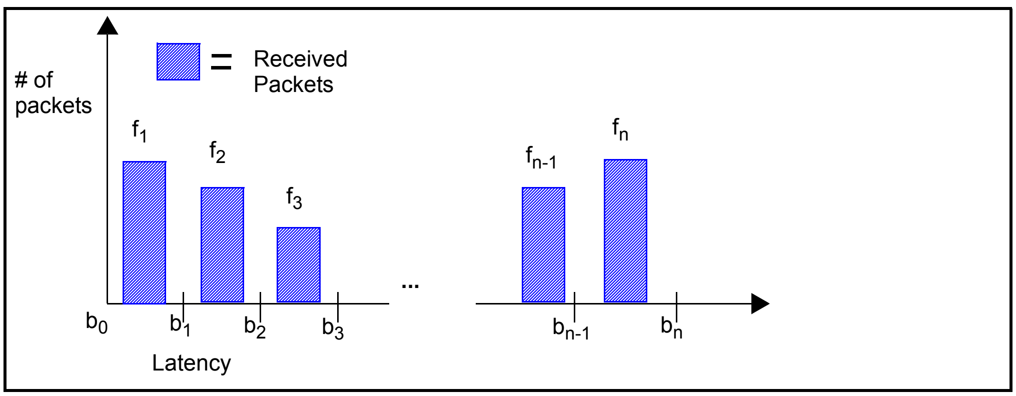

Figure 13-1. Latency Distribution Buckets with Packets Received with Various Latencies

1 The probability mass function is calculated as:

$$p(x_i) = P\{X = x_i\}$$

where $P\{X = x_i\}$ is the probability that $X = x_i$.

The probability of $X = x_i$ is calculated as::

$$P\{X = x_i\} = \frac{f_i}{\sum_{k=1}^{n} f_k}$$

where k is another integer (such as 1, 2, 3) other than i.

2 The expectation of random variable X is calculated as:

$$E[X] = \sum_{i=1}^{n} x_i \; p(x_i)$$

where $p(x_i)$ is the probability mass function defined in *Step 1*.

$E[X]$ is also known as the first moment of random variable X.

3 The second moment of random variable X is calculated as:

$$E[X^2] = \sum_{i=1}^{n} x_i^2 \; p(x_i)$$

4 The variance of random variable X is calculated as:

$$VAR[X] = (E[X^2] - (E[X])^2)$$

5 The standard deviation of random variable X is calculated as:

$$STD[X] = \sqrt{VAR[X]}$$

Standard deviation is also commonly denoted by σ.

In SmartFlow, this is the standard deviation of the latency distribution.

Setting Latency Intervals

Specify the time intervals (latency buckets) into which the latency of each group and/or flow should fall for the Latency Distribution test in the *Latency Distribution* pane on the *Test Setup>Individual Tests* tab. *Figure 13-2* shows the default time intervals.

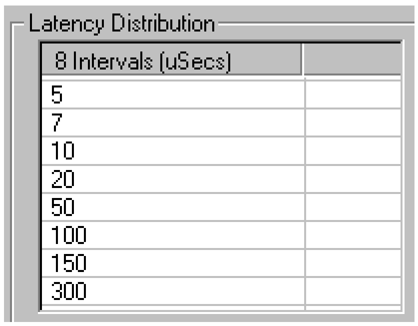

Figure 13-2. Latency Distribution Pane on Individual Tests Tab

The upper boundary of a latency bucket is the lower boundary of the next bucket. Thus, if two of the bucket intervals were 100 and 500, packets whose latency fell into the 500 bucket in results would range anywhere from 101 to 499 microseconds.

The maximum value that can be set for the largest bucket is 429,496,728. The eighth bucket acts as a catch all for frames that do not fall into the first seven buckets.

Latency Interval Guidelines

If devices that are slow (higher latency) are being tested, set the latency interval values to higher numbers.

- *For switches:*
  Since switches are normally fast, set the latency buckets to low values.
- *For routers:*
  Since routers are normally slower than switches, set the latency buckets to higher values.
- *For wireless devices:*
  Set the latency buckets to higher values than would be set for a router.

Once the test is run, you may want to see a more detailed breakdown of a latency interval or determine exactly into what range latency begins. To do this, find the peak(s) within the latency buckets and create new buckets around these values.

Example: In this example, *Figure 13-3* shows a Latency Distribution chart that contains two peaks: one in the 10 bucket and one in the 1,000 bucket.

In this case, you might create another test with four latency buckets around the 0 through 10 peak and another four buckets around the 1,000 peak.

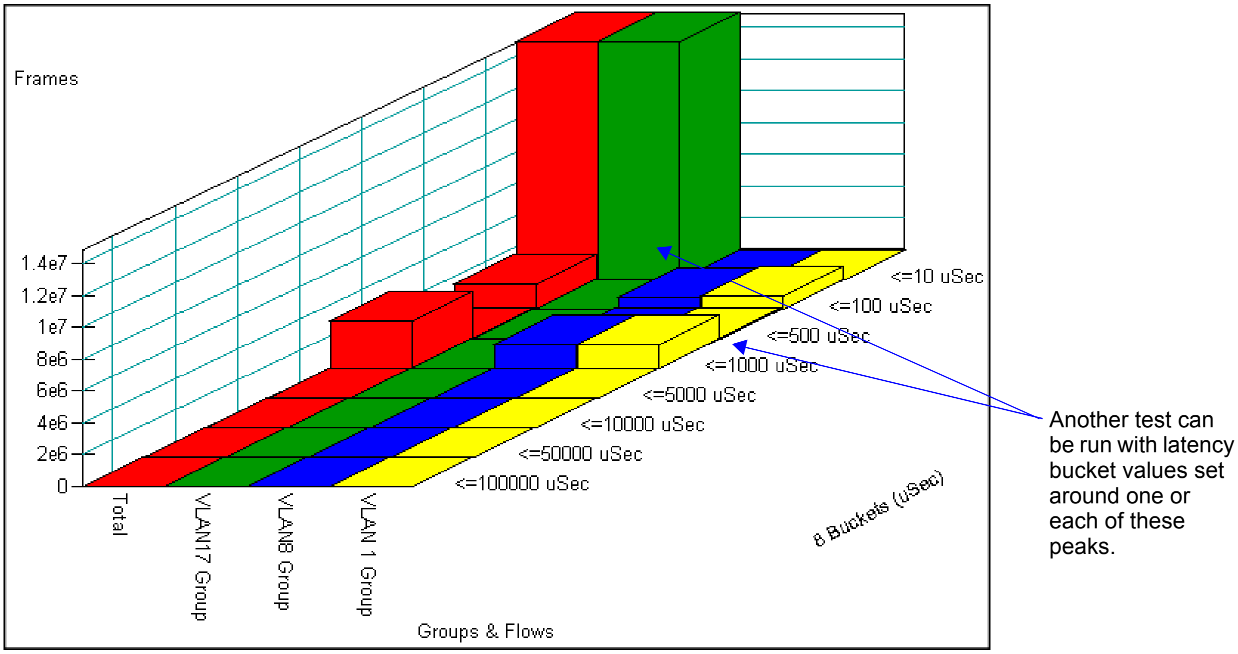

Figure 13-3. Latency Distribution Chart Example

Interpreting Latency Distribution Test Results

Results of the Latency Distribution test appear in the main window at the conclusion of the test. View test results for this test in graphical or tabular form. A summary or detailed report can be viewed in tabular form.

The standard deviation for the Latency Distribution test depends in part on the size of the latency buckets that were specified in the test setup. If many buckets (intervals) were defined that vary only slightly in size from one another, the distribution spans multiple buckets and the standard deviation is larger than if larger buckets were specified. If latency falls all into one bucket, the standard deviation is zero (0).

Notes:

- If *all* of the flows and/or groups appear to fall into the largest latency bucket, then the counters might be overflowed. On the *Test Setup>Individual Tests* tab, change the last latency bucket to a larger number to handle slower devices and then rerun the test. (The maximum value for the largest bucket is 429,496,728.)
- Results formats can vary depending on whether or not frame size automation was used during the test. For more information, see *“Using Frame Size Automation (Global for All Flows)” on page 236*.

Chart (Latency Distribution)

The Latency Distribution chart results graph the pattern of latency across eight variable time buckets. If frame size automation was used, both the frame size as well as the test iteration can be selected at the top of the results tab.

In *Figure 13-4*, all latency values fell between 11 and 599 microseconds at test iteration 8.

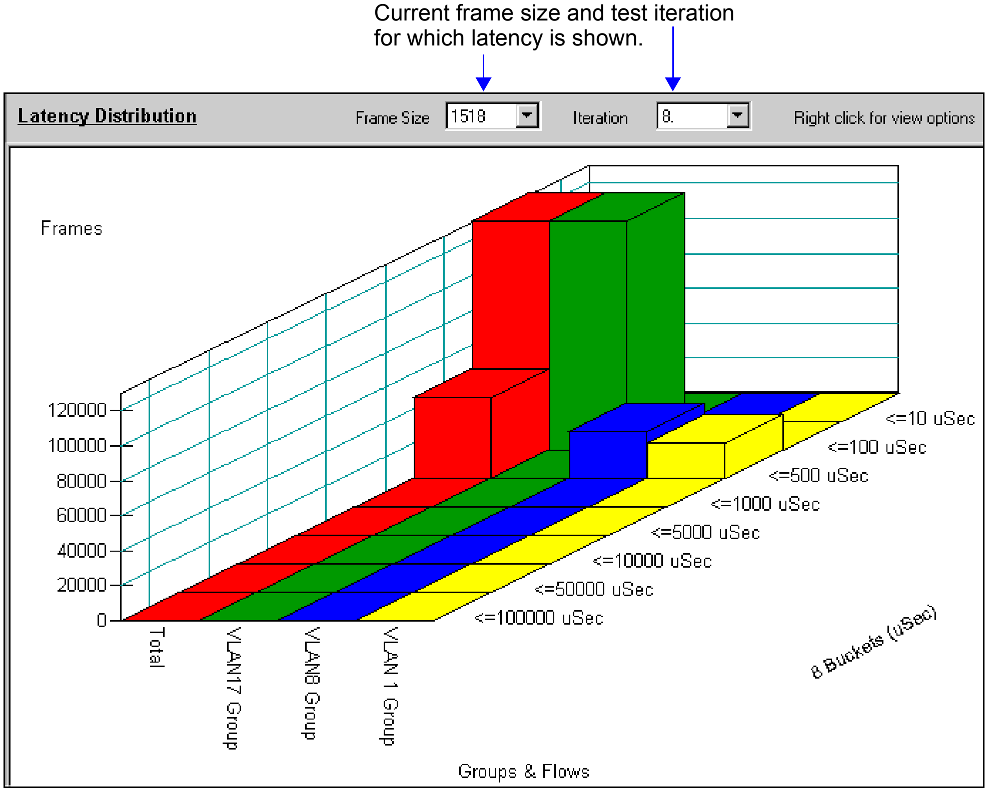

Figure 13-4. Latency Distribution Chart Results

The following table lists the fields in the Latency Distribution chart results.

For a description of each field in Latency Distribution chart results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Summary Table (Latency Distribution)

The Latency Distribution summary results table shows (by traffic load) how the latency for each group (or flow) falls across eight variable time interval buckets. (See *Figure 13-5*.) If frame size automation was used, the frame size can be selected at the top of the results tab.

The number of frames peaks with latency between 11 and 100 microseconds, and again at the 1,001- to 5,000-microsecond range.

Name/Bucket	Framesi	Load (%)	<=10 uSec	<=100 uSec	<=500 uSec	<=1000 uSec	<=5000 uSec	<=10000 uS
Total	64	100.00	0	2976191	208	2976127	0	0
Vlan1	64	100.00	0	26	119	1488032	0	0
Vlan 2	64	100.00	0	1488072	0	0	0	0
Vlan 3	64	100.00	0	21	89	1488095	0	0
Vlan 4	64	100.00	0	1488072	0	0	0	0

Figure 13-5. Latency Distribution Summary Results

Check if the device is giving priority to groups with higher priority. The latency that does occur for higher priority groups should be in the lower latency ranges, whereas the lower priority groups have more frames that fall into the higher time intervals.

For a description of each field in Latency Distribution summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Detailed Table (Latency Distribution)

The Latency Distribution detailed results table shows (by load percentage) the latency distribution for each group and/or flow in eight variable time buckets as well as latency standard deviation for the group and/or flow. (See *Figure 13-6*.) Totals for the load are displayed at the start of each new load percentage.

If, for example, a bucket for <=10 and another for <=100 was specified, all frames with a latency of 10 microseconds fall into the first <=10 bucket.

Total for all groups/flows at 10% load →

Latency values for next load (20) percent starts here. →

← More groups/ flows had latency in this interval once the load was increased.

Name	ILoad	StdDeviati	TxFrames	RxFrames	<=5 uSec	<=7 uSec	<=20 uSec
Total	10.000	0.00	101,344	101,344	0	0	101,344
No Vlan Group	10.000	0.00	50,672	50,672	0	0	50,672
Vlans Group	10.000	0.00	50,672	50,672	0	0	50,672
No Vlan 1A1-0->1	5.0000	0.00	12,668	12,668	0	0	12,668
No Vlan 1A2-0->1	5.0000	0.00	12,668	12,668	0	0	12,668
No Vlan 6A1-0->6	5.0000	0.00	12,668	12,668	0	0	12,668
No Vlan 6A2-0->6	5.0000	0.00	12,668	12,668	0	0	12,668
Vlans 1A1-1-v101-	5.0000	0.00	12,668	12,668	0	0	12,668
Vlans 1A2-1-v102-	5.0000	0.00	12,668	12,668	0	0	12,668
Vlans 6A1-1-v111-	5.0000	0.00	12,668	12,668	0	0	12,668
Vlans 6A2-1-v112-	5.0000	0.00	12,668	12,668	0	0	12,668
Total	20.000	0.00	202,696	202,696	0	0	202,696
No Vlan Group	20.000	0.00	101,348	101,348	0	0	101,348
Vlans Group	20.000	0.00	101,348	101,348	0	0	101,348
No Vlan 1A1-0->1	10.000	0.00	25,337	25,337	0	0	25,337
No Vlan 1A2-0->1	10.000	0.00	25,337	25,337	0	0	25,337
No Vlan 6A1-0->6	10.000	0.00	25,337	25,337	0	0	25,337
No Vlan 6A2-0->6	10.000	0.00	25,337	25,337	0	0	25,337
Vlans 1A1-1-v101-	10.000	0.00	25,337	25,337	0	0	25,337
Vlans 1A2-1-v102-	10.000	0.00	25,337	25,337	0	0	25,337
Vlans 6A1-1-v111-	10.000	0.00	25,337	25,337	0	0	25,337
Vlans 6A2-1-v112-	10.000	0.00	25,337	25,337	0	0	25,337

Figure 13-6. Latency Distribution Detailed Results

Note: When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

For a description of each field in the Latency Distribution test detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Stray Frames Results Table (Latency Distribution)

The Latency Distribution stray frames results table shows (per port) frames that were received but were not destined (expected) for that port. (See *Figure 13-7*.) It also lists each flow (not SmartFlow) containing the stray frames received by that port. Stray frames are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Stray frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* results tab even if other flows in the same SmartFlow are listed.

Latency Distribution

Time	ILoad	FrameS	PortName	FlowName	SourceIP	DestIP	RxFram	StdDev	<=5 uSec	<=7 uSec
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow NoVLAN 6A1-0->6A2-0	133.11.0.3	133.12.0.3	12476	0	12476	0
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow NoVLAN 6A2-0->6A1-0	133.12.0.3	133.11.0.3	9510	0	9510	0
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow VLAN 6A1-1-v111->6A2-1-v112	30.111.0.3	30.112.0.3	12475	0	12475	0
8/8/03 3:30:45 PM	10	128	Port SF 1A1	Flow VLAN 6A2-1-v112->6A1-1-v111	30.112.0.3	30.111.0.3	9509	0	9509	0
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow NoVLAN 6A1-0->6A2-0	133.11.0.3	133.12.0.3	12476	0	12476	0
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow NoVLAN 6A2-0->6A1-0	133.12.0.3	133.11.0.3	9510	0	9510	0
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow VLAN 6A1-1-v111->6A2-1-v112	30.111.0.3	30.112.0.3	12475	0	12475	0
8/8/03 3:30:45 PM	10	128	Port SF 1A2	Flow VLAN 6A2-1-v112->6A1-1-v111	30.112.0.3	30.111.0.3	9509	0	9509	0
8/8/03 3:30:45 PM	10	128	Port SF 6A1	Flow NoVLAN 1A2-0->1A1-0	133.2.0.3	133.1.0.3	9369	0	9368	0
8/8/03 3:30:45 PM	10	128	Port SF 6A1	Flow NoVLAN 1A1-0->1A2-0	133.1.0.3	133.2.0.3	9369	0	9368	0

Figure 13-7. Latency Distribution Stray Frames Tab

Each row in the *Stray Frames* results tab represents one flow of a SmartFlow. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

Number of Records in Stray Frames Results

When *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. You can adjust this default limit by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

For a description of each field in Latency Distribution test stray frame results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

SPIRENT
Communications

Chapter 14

Latency Over Time Test

In this chapter...

About the Latency Over Time Test

Use the Latency Over Time test to determine the buffering capabilities of a device and to see whether the algorithm being used for buffering is adequate. The Latency Over Time test shows the overall pattern of latency across time as the test is run at different loads.

For each receiving port, Latency Over Time results show the minimum, average, and maximum latency of frames received during the specified time interval throughout the test. For example, if the interval specified is 5 seconds and the test duration is 20 seconds, results display the latency that exists for frames at 5, 10, 15, and 20 seconds into the test. For all TeraMetrics and TeraMetrics-based receiving ports, it calculates either average and maximum latency or minimum and maximum latency of frames received during the specified time interval throughout the test.

Compared to the Latency test, the Latency Over Time test can provide a more detailed view of latency behavior at the DUT load tolerance limits. Whereas the Latency Distribution test shows where latency falls according to *amount of latency*, the Latency Over Time test shows at what point(s) during the test that the latency occurs.

Note: WAN cards do not support the Latency Over Time test.

Test Methodology

This test tracks the latency of each test frame on a time interval basis and places latency results into a column that represents one time period at which measurements are taken. The time period is based on the value specified in the *Time interval (Sec)* field on the *Test Setup>Individual Tests* tab. SmartFlow reads the time that the sending SmartBits port sent the frame (transmit timestamp) and the time that the receiving SmartBits recognizes the trigger frame, which is the receive timestamp.

Latency is calculated as:

Latency = (Receive Timestamp) - (Transmit Timestamp)

You can set up multiple unicast flows for each port, some unidirectional and some bi-directional, each with a different frame size, protocol, and scheme for priority routing.

The test uses the specified test duration; a starting percentage load (based on the wire rate); and either custom loads, or a step percentage by which to increase the load during the test and a stop percentage load at which the test ends. During the test, these values determine the duration for which the DUT is tested at a specific load.

Note: Latency values are only for the frames that were not dropped.

Setting a Time Interval for Latency Over Time

The value in the *Time interval (Sec)* field on the *Test Setup>Individual Tests* tab determines how often SmartFlow takes measurements during a test and whether or not the results are broad or focused.

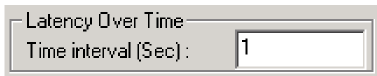

The maximum value that can be set for the interval is less than or equal to the value set in the *Duration Time (Sec.)* field on the *Test Setup>Test Iterations* tab. For example, if the test duration is 100, the interval cannot be greater than 100. Smaller intervals result in more focused results. The larger the interval, the fewer measurements occur.

Tip: For a more granular, focused test, set the interval value to 1 (minimum).

For example, if an interval of 10 seconds is specified for a 20-second test, the trend for the entire test is evident. If an interval of 1 second is specified, 20 measurements are taken and 20 columns appear in results.

Interpreting Latency Over Time Test Results

Results of the Latency Over Time test are based on receiving ports. Results appear in the main window at the conclusion of the test. Test results for this test can be viewed in graphical or tabular form. A summary, detailed, or specialized (such as Stray Frames) report can be viewed in tabular form.

An ideal DUT would slightly increase latency for a short time. Increasing latency indicates that the DUT has started buffering more and more traffic. If the latency increases from start to finish during a test, it indicates that the DUT cannot handle the load for sustained periods of time and most likely frame loss begins or increases. Latency that is rising steeply (a spike) indicates that the device buffer has overflowed and there is a problem with buffer scheduling. If there is a steady, gradual increase in latency, it is likely that the DUT is able to keep up with the inflow of frames.

Note: Results formats can vary depending on whether or not varying frame sizes or test loads were used for the test. For more information, see *"Varying Frame Sizes" on page 225* or *"Varying Test Loads" on page 253*.

Chart (Latency Over Time)

The Latency Over Time chart results graph a pattern of latency across time as the test is run (*Figure 14-1*). For each receiving port, it shows the average latency in milliseconds at regular intervals throughout the test.

If frame size automation was used, both the frame size as well as the test iteration can be selected at the top of the results tab.

Figure 14-1 shows a steady but gradual increase in latency, which can indicate the DUT is handling the frames that it receives with buffering. Each group has almost the same latency, regardless of the priority of the flows.

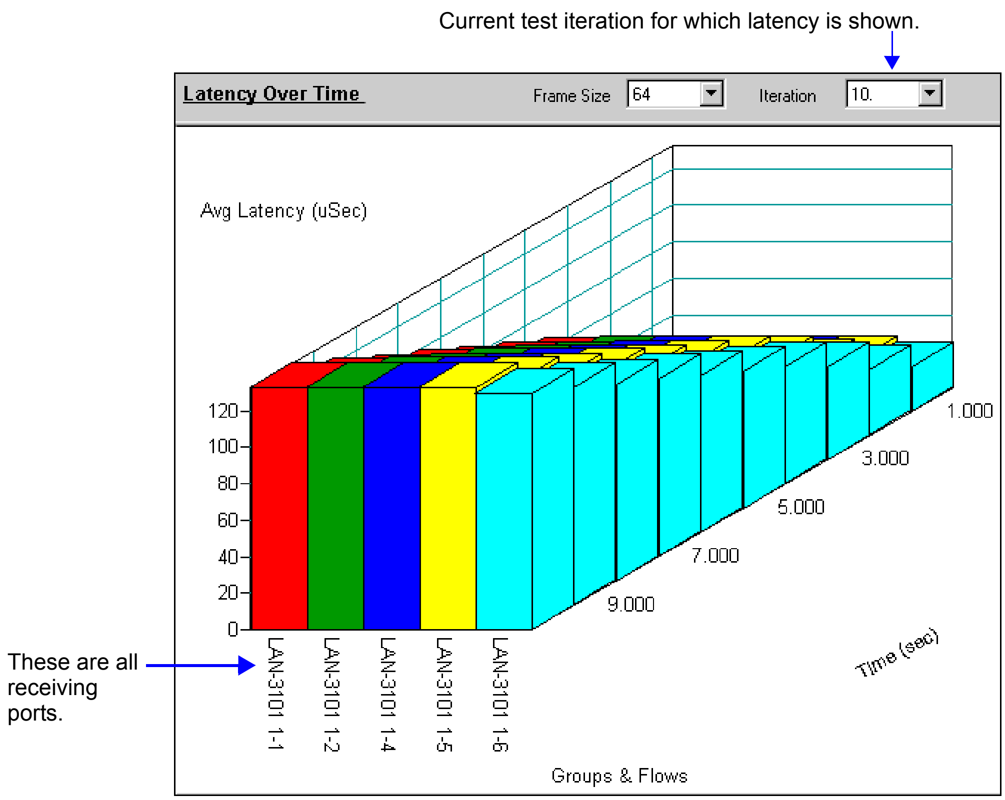

Figure 14-1. Latency Over Time Chart Results

For a description of each field in Latency Over Time test chart results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Summary Table (Latency Over Time)

The Latency Over Time summary results table shows where latency occurs at different points throughout the test. It shows the pattern of latency (for each receiving port) at regular intervals during each iteration of the test.

Select the test iteration for which you want to see latency at the top of the results tab. (See *Figure 14-2*.) If frame size automation was used, you can also select the frame size at the top of the results tab.

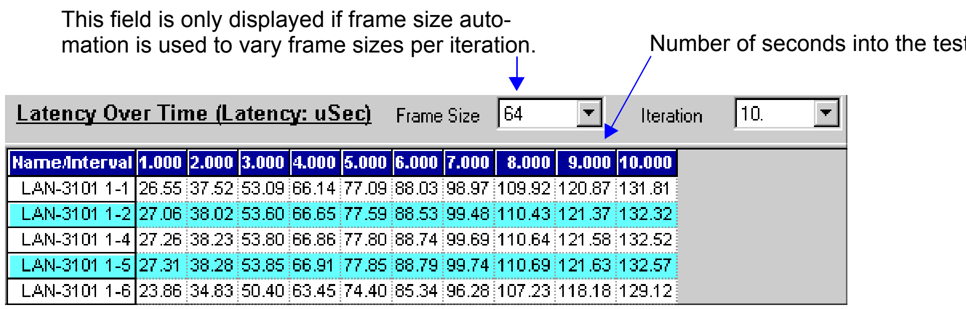

Name/Interval	1.000	2.000	3.000	4.000	5.000	6.000	7.000	8.000	9.000	10.000
LAN-3101 1-1	26.55	37.52	53.09	66.14	77.09	88.03	98.97	109.92	120.87	131.81
LAN-3101 1-2	27.06	38.02	53.60	66.65	77.59	88.53	99.48	110.43	121.37	132.32
LAN-3101 1-4	27.26	38.23	53.80	66.86	77.80	88.74	99.69	110.64	121.58	132.52
LAN-3101 1-5	27.31	38.28	53.85	66.91	77.85	88.79	99.74	110.69	121.63	132.57
LAN-3101 1-6	23.86	34.83	50.40	63.45	74.40	85.34	96.28	107.23	118.18	129.12

Figure 14-2. Latency Over Time Summary Results

In this example, as the test progressed, latency steadily increased for all flows. The device did not give much priority to any flows except the 1-6.

For a description of each field in Latency Over Time test summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Detailed Table (Latency Over Time)

For each receiving port, the Latency Over Time detailed results include the minimum, average, and maximum latency at each specified interval of a test iteration. (See *Figure 14-3*.) If you are using frame size automation to vary frame sizes, results are grouped first by frame size and then by receiving port.

For example, if the test included frames at sizes 64 and 128, results would first display latency for all ports receiving frames at 64 bytes, then all flows for ports receiving frames at 128 bytes.

Note: For tests using custom frame sizes per flow or custom loads, N/A appears in the *FrameSize* and *Load* columns respectively. For more information about custom frame sizes or loads, see *"Using Custom Frame Sizes per Flow" on page 228* or *"Varying Test Loads" on page 253*.

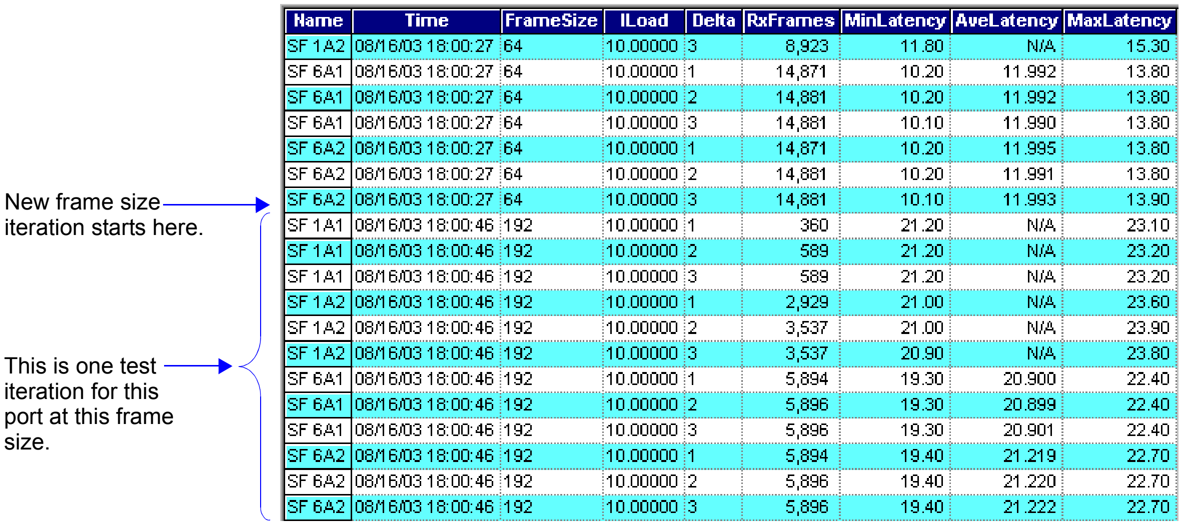

Name	Time	FrameSize	ILoad	Delta	RxFrames	MinLatency	AveLatency	MaxLatency
SF 1A2	08/16/03 18:00:27	64	10.00000	3	8,923	11.80	N/A	15.30
SF 6A1	08/16/03 18:00:27	64	10.00000	1	14,871	10.20	11.992	13.80
SF 6A1	08/16/03 18:00:27	64	10.00000	2	14,881	10.20	11.992	13.80
SF 6A1	08/16/03 18:00:27	64	10.00000	3	14,881	10.10	11.990	13.80
SF 6A2	08/16/03 18:00:27	64	10.00000	1	14,871	10.20	11.995	13.80
SF 6A2	08/16/03 18:00:27	64	10.00000	2	14,881	10.20	11.991	13.80
SF 6A2	08/16/03 18:00:27	64	10.00000	3	14,881	10.10	11.993	13.90
SF 1A1	08/16/03 18:00:46	192	10.00000	1	360	21.20	N/A	23.10
SF 1A1	08/16/03 18:00:46	192	10.00000	2	589	21.20	N/A	23.20
SF 1A1	08/16/03 18:00:46	192	10.00000	3	589	21.20	N/A	23.20
SF 1A2	08/16/03 18:00:46	192	10.00000	1	2,929	21.00	N/A	23.60
SF 1A2	08/16/03 18:00:46	192	10.00000	2	3,537	21.00	N/A	23.90
SF 1A2	08/16/03 18:00:46	192	10.00000	3	3,537	20.90	N/A	23.80
SF 6A1	08/16/03 18:00:46	192	10.00000	1	5,894	19.30	20.900	22.40
SF 6A1	08/16/03 18:00:46	192	10.00000	2	5,896	19.30	20.899	22.40
SF 6A1	08/16/03 18:00:46	192	10.00000	3	5,896	19.30	20.901	22.40
SF 6A2	08/16/03 18:00:46	192	10.00000	1	5,894	19.40	21.219	22.70
SF 6A2	08/16/03 18:00:46	192	10.00000	2	5,896	19.40	21.220	22.70
SF 6A2	08/16/03 18:00:46	192	10.00000	3	5,896	19.40	21.222	22.70

Figure 14-3. Latency Over Time Detailed Results

Note: When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

For a description of each field in the Latency Over Time test detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Stray Frames Results Table (Latency Over Time)

The Latency Over Time test does not report stray frames results.

Chapter 15

Latency SnapShot Test

In this chapter...

Latency SnapShot Test

After determining load tolerance limits with the Frame Loss test, the Latency SnapShot test can be used to examine latency in a sample consisting of a user-defined number of frames. However, you can also run the Latency SnapShot test any time without running any prior test. Compared to the Latency and Latency Distribution tests, this test can provide the most magnified view of latency behavior at the DUT load tolerance limits, allowing you to precisely analyze jitter implied in results for the other tests.

Test Methodology

This test tracks the latency of each test frame on a frame-by-frame basis and then captures the specified number of packet tags for the snapshot. SmartFlow reads the time that the sending SmartBits port sent the frame (transmit timestamp) and the time that the receiving SmartBits recognizes the trigger frame, which is the receive timestamp.

Latency is calculated as:

(Receive Timestamp) - (Transmit Timestamp) = Latency

You can set up multiple unicast flows for each port, some unidirectional and some bi-directional, each with a different frame size, protocol, and scheme for priority routing.

The test uses the specified test duration, a starting percentage load (based on the wire rate), a step percentage by which to increase the load during the test, and a stop percentage load at which the test ends. During the test, these three values determine the duration for which the DUT is tested at a specific load.

Note: Latency values are only for the frames that were not dropped.

Interpreting Latency SnapShot Test Results

Results of the Latency SnapShot test appear in the main window at the conclusion of the test. Test results for this test can be viewed in graphical or tabular form. A summary or detailed report can be viewed in tabular form.

Results format can vary depending on whether or not frame size automation was used during the test. See *"Using Frame Size Automation (Global for All Flows)" on page 236*.

Chart (Latency SnapShot)

For flows and groups involved in a Latency SnapShot test, a graph of results visualizes a relative pattern of jitter within the sample of frames. (See *Figure 15-1*.) If frame size automation was used, both the frame size as well as the test iteration can be selected at the top of the results tab.

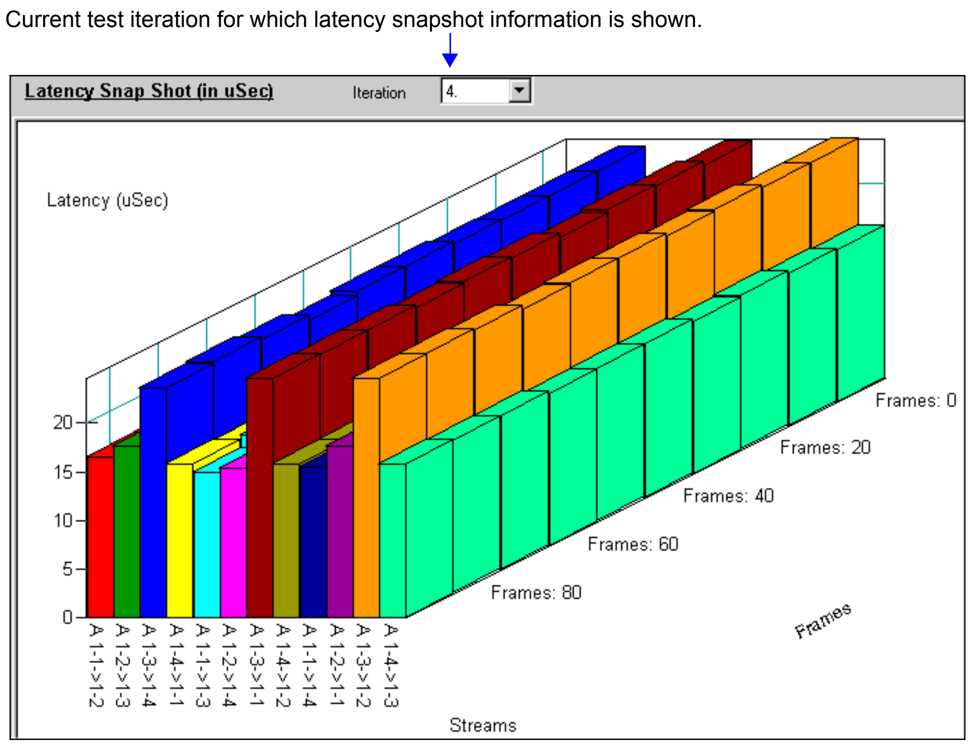

Figure 15-1. Latency SnapShot Chart Results

For a description of each field in Latency SnapShot test chart results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Summary Table (Latency SnapShot)

For flows and groups involved in a Latency SnapShot test that does not use frame size automation or the *Custom frame sizes (per flow)* option on the *Test Setup>Test Iterations* tab, the summary table of results lists the latency for each group of 10 frames in the sample. (See *Figure 15-2*.)

Latency Snap Shot (in uSec) Iteration 4.

Name/Frame	Frames: 0	Frames: 10	Frames: 20	Frames: 30	Frames: 40	Frames: 50
A 1-1->1-2	16.40	16.30	16.40	16.30	16.40	16.30
A 1-2->1-3	17.30	15.30	17.20	17.40	17.50	17.50
A 1-3->1-4	23.70	23.60	23.40	23.50	23.70	23.50
A 1-4->1-1	15.30	15.60	15.20	15.70	15.80	15.70
A 1-1->1-3	16.20	15.00	16.40	16.30	16.40	15.20
A 1-2->1-4	16.70	15.40	16.50	16.60	16.60	15.10
A 1-3->1-1	24.20	24.50	24.60	24.30	24.50	24.50
A 1-4->1-2	15.60	15.60	15.70	15.60	15.70	15.70
A 1-1->1-4	15.50	15.50	15.40	15.40	15.60	15.40
A 1-2->1-1	17.20	17.50	17.50	17.20	17.60	17.50
A 1-3->1-2	24.50	24.40	24.50	24.40	24.50	24.40
A 1-4->1-3	15.50	15.60	15.70	15.60	15.70	15.60

Figure 15-2. Latency SnapShot Summary Results without Frame Size Automation

If the test is run using frame size automation, the Latency Snapshot test results provide a view for each rate of each frame size. (See *Figure 15-3*.) Both the frame size as well as the load can be selected at the top of the results tab.

Select the frame size and load to view.

Latency Snap Shot (in uSec) Frame Size 128 Iteration 3.

Name/Frame	Frames: 0	Frames: 10	Frames: 20	Frames: 30	Frames: 40	Frames: 50	Frames: 60
A 1-1->1-2	17.60	17.10	17.60	17.50	17.60	15.30	17.50
A 1-2->1-3	17.20	16.80	17.30	17.30	17.20	15.40	17.10
A 1-3->1-4	15.50	15.10	15.50	15.30	15.50	15.40	15.30
A 1-4->1-1	15.70	15.60	15.70	15.60	15.70	15.70	15.60
A 1-1->1-3	17.50	17.10	17.50	17.70	17.60	15.70	17.40
A 1-2->1-4	16.40	16.50	16.40	16.30	16.40	16.30	16.10
A 1-3->1-1	16.30	16.30	16.40	16.20	16.40	16.30	16.20
A 1-4->1-2	15.70	15.60	15.70	15.60	15.70	15.70	15.50
A 1-1->1-4	16.70	16.80	16.80	16.70	16.70	16.70	16.40
A 1-2->1-1	17.30	17.00	17.30	17.20	17.30	17.30	17.20
A 1-3->1-2	16.30	15.90	16.30	16.20	16.40	16.20	16.20
A 1-4->1-3	15.70	15.50	15.60	15.70	15.60	15.20	15.70

Figure 15-3. Latency SnapShot Summary Results with Frame Size Automation

For a description of each field in Latency SnapShot test summary results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Detailed Table (Latency SnapShot)

For flows and groups involved in a Latency SnapShot test at a given load, a detailed table lists the latency for each frame in a user-defined 100-frame sample. (See *Figure 15-4*.)

Stream	Load (%)	Frame: 0 (uSecs)	Frame: 1 (uSecs)	Frame: 2 (uSecs)	Frame: 3 (uSecs)	Frame: 4 (uSecs)	Frame: 5 (uSecs)
A 1->2	10.00000	0.2	0.1	0.1	0.1	0.1	0.1
A 2->3	10.00000	0.1	0.1	0.1	0.1	0.1	0.1
A 3->4	10.00000	0.2	0.2	0.2	0.2	0.2	0.2
A 4->1	10.00000	0.1	0.1	0.1	0.1	0.1	0.1
A 1->3	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 2->4	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 3->1	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 4->2	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 1->4	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 2->1	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 3->2	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0
A 4->3	10.00000	1933313.2	2038523.6	0.0	0.0	0.0	0.0

Figure 15-4. Latency SnapShot Detailed Results

Note: When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

For a description of each field in the Latency SnapShot test detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Stray Frames Results Table (Latency SnapShot)

The Latency SnapShot stray frames results table shows (per port) frames in the snapshot that were received but were not destined (expected) for that port. (See *Figure 15-5*.) It also lists each flow (not SmartFlow) in the snapshot containing the stray frames received by that port. Stray frames are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Stray frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* results tab even if other flows in the same SmartFlow are listed.

Time	FrameSize	ILoad	PortName	FlowName	SourceIP	DestIP	RxFrames	DeltaID
12/4/01 7:26:29 PM	128	50	Port SMB 21 4	Flow A 1->2-0	192.85.1.3	192.85.2.3	100	0
12/4/01 7:26:29 PM	128	50	Port SMB 21 4	Flow A 1->2-1	192.85.1.3	192.85.2.4	100	0
12/4/01 7:26:29 PM	128	50	Port SMB 21 4	Flow A 1->2-2	192.85.1.3	192.85.2.5	100	0
12/4/01 7:26:29 PM	128	50	Port SMB 21 4	Flow A 1->2-3	192.85.1.3	192.85.2.6	100	0
12/4/01 7:26:29 PM	128	50	Port SMB 21 4	Flow A 1->2-4	192.85.1.3	192.85.2.7	100	0
12/4/01 7:26:45 PM	128	100	Port SMB 21 4	Flow A 1->2-0	192.85.1.3	192.85.2.3	100	0
12/4/01 7:26:45 PM	128	100	Port SMB 21 4	Flow A 1->2-1	192.85.1.3	192.85.2.4	100	0
12/4/01 7:26:45 PM	128	100	Port SMB 21 4	Flow A 1->2-2	192.85.1.3	192.85.2.5	100	0
12/4/01 7:26:45 PM	128	100	Port SMB 21 4	Flow A 1->2-3	192.85.1.3	192.85.2.6	100	0
12/4/01 7:26:45 PM	128	100	Port SMB 21 4	Flow A 1->2-4	192.85.1.3	192.85.2.7	100	0

Figure 15-5. Latency SnapShot Stray Frames Results

Each row in the *Stray Frames* results tab represents one flow of a SmartFlow. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

Number of Records in Stray Frames Results

When *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. You can adjust this default limit by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

For a description of each field in Latency SnapShot test stray frames results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Chapter 16

SmartTracker Test

In this chapter...

About the SmartTracker Test

The SmartTracker test enables you to track a specified field in test packets and obtain measurements based on the tracked field. It provides predefined fields to track, as well as the flexibility for custom tracking. Trackable fields include:

- QoS (predefined) – ToS and DiffServ IP precedence
- VLANs – tagging, priority, and leakage
- Cyclic flows
- Any other custom, user-defined field – for example, MPLS label, TCP, or UDP destination port.

SmartTracker uses TeraMetrics and TeraMetrics-based modules exclusively, with their *Alternate Key* feature, which enables you to focus tracking on a particular field in a packet. SmartTracker supports both unicast and multicast flows as well as IPv4 and IPv6 protocols.

Important: SmartTracker will track an IPv6 frame with an extension header but will not track the portion of the frame that is the IPv6 extension header.

Note: SmartTracker detects and tracks only frames that contain the SmartBits signature field. Miscellaneous frames, such as management frames or Spanning Tree bridge protocol data units (BPDUs), are not included in results.

SmartTracker results provide not only received statistics for a particular field, but a correlation between the received and transmitted values to ensure that you are comparing received values for the correct flow. If you want to track received statistics for a large number of flows without any correlation to the transmitted value, you can use another SmartFlow test, and you can track up to 65,536 cyclic flows per port with standard tracking based on the SmartBits signature field.

The SmartTracker test provides the ability to test:

- **QoS and a DUT's ability to properly reclassify packets according to IP precedence settings.**
  A popular QoS technique is to set the IP precedence bits in the ToS and DiffServ fields of packets. Some DUTs are designed to change the value of the IP precedence field based on certain packet characteristics, such as protocol type. The SmartTracker test provides a way to track flows based on the IP precedence settings and perform measurements based on this type or reclassification. It can track the TOS and DiffServ fields in IPv4 flows.
  This test can answer the question "*Is my router classifying QoS traffic correctly*?"
- **VLANs, VLAN tags, VLAN leakage, and VLAN stacking (QinQ).**
  The SmartTracker test provides metrics for received frames that meet certain criteria relative to the VLAN tag of the frame. It can measure VLAN leakage of traffic across VLAN boundaries, which can compromise security. The test allows you to track other

criteria relative to a frame VLAN tag. For example, you can test VLAN priority or determine whether the VLAN tag was retained in received frames and whether it was in the correct place. For more information, see *"VLAN Testing" on page 393*.

This test can answer the question "*Is my VLAN traffic staying within the right VLAN?*"

- **Cyclic flows.**
  Testing with cyclic SmartFlows (as opposed to multiple non-cyclic SmartFlows) increases the number of flows that can be transmitted and tracked. One cyclic flow uses only one hardware stream regardless of the number of field variations, whereas each individual SmartFlow requires a separate hardware stream. For example, if the test is set up to track eight bits of the TCP destination port field, the test can track 255 possible values of the TCP port with just one flow.

  SmartTracker is the only test that enables you to track cyclic flows individually in results, rather than collectively as a group of cyclic flows (one SmartFlow). By using the *Custom (user-defined)* option in the *Tracking and Filter Options* field on the *Test Setup>SmartTracker* tab, you can track a portion of the source or destination IP address. By tracking each variation of the source or destination IP address, the cyclic flow is actually being tracked. For more information, see *"How Cyclic Flows are Tracked" on page 393*.

- **Custom tracking of a user-defined field.**

  SmartTracker enables you to track any field within an IP packet and label the results column in which the information appears. Custom tracking limits the number of trackable flows. For more information, see *"Limitations" on page 388*.

  Other entities that can be tracked are MPLS label; TCP destination port; UDP destination port; and other user-defined tracking, by specifying a given offset and size up to 16 bits, depending on what needs to be tracked.

Tracking and Filtering Options

The SmartTracker test offers predefined or custom tracking, depending on what needs to be tested:

- Tracking the QoS DiffServ or ToS IP precedence field in the packet uses a predefined tracking setup.
- Tracking another (custom) field in the packet requires a custom setup.
  This includes any field other than the stream ID field, such as a VLAN tag, TCP destination port, or any other custom field.

Traffic filters that are specified in the test setup act as pre-capture filters and filter traffic out (or in) prior to tracking for QoS or a specified field and collecting results. For example, selecting the *QoS (IP Priority) on IPv4 with VLAN* option in the *Tracking and Filter Options* pane on the *Test Setup>SmartTracker* tab causes SmartFlow to examine packets for the presence of an IPv4 version number and a 0x8100 value in the *Type* field (which indicate the presence of a VLAN tag). Only traffic that meets this criteria is included for QoS tracking in results.

QoS Tracking and Predefined Filters

Each QoS tracking option includes a predefined, inclusive traffic filter. SmartFlow only examines the type of traffic that includes the specified criteria (field). For example, if the *QoS (IP Priority) on IPv4 with VLAN* option is selected in the *Tracking and Filter Options* pane, results provide information about the QoS field only for IPv4 flows with VLAN tags.

Notes:
- The offset values displayed for the predefined filters are based on Ethernet frames. If a transmitting or receiving port uses a different technology (e.g., POS), SmartFlow automatically adjusts the offset based on the location of the field that is being tracked in an Ethernet frame.
- Predefined filters with VLAN or MPLS only include frames with a single VLAN or MPLS tag. Results for packets with multiple tags produce unreliable results. To track flows that contain multiple VLAN or MPLS tags, use the *Custom (user defined)* option in the *Tracking and Filter Options* pane.

Custom Tracking and Filters

Custom tracking allows you to set tracking criteria by specifying the location of the field (size in bits and offset) to track. For custom tracking, it is recommended that a predefined inclusive traffic filter is selected or a custom filter is defined that can include or exclude traffic types. For example, to track the TCP port field, include a filter field value that contains *IPv4* so that any non-IPv4 flows (including IPv4 with tags) are not included in results.

DSCP Tracking

To track all 6 bits of the DSCP, in the *Tracking and Filter Options* pane, use *Custom (user defined)* as the field value.

In the *Offset (bytes)* field, enter a value of 15 and select 6 bits to track (since DSCP is 6 bits long) in the *Bits to track* field. [To include the explicit congestion control bits (2 bits), enter a value of 8 in the *Bits to track* field.]

Important: These settings reflect the fact that no VLAN tags are present in the packets. If there are VLAN tags in the expected traffic, then set the offset to 17 bytes (since the VLAN tag is 2 bytes). Optionally, if you expect to receive a mixture of tagged and untagged VLANS in the traffic, filter out one or the other using the fields in the *Traffic Filter* pane.

When defining a custom setup, the number of SmartFlows that can be tracked in the test is determined by the number of bits that is being tracked.

Tip: For custom tracking, it is recommended that you select a predefined inclusive traffic filter or define a custom filter than can include or exclude traffic types. For example, to track the TCP port field, include a filter of *IPv4* in the *Traffic Filter* pane so that any non-IPv4 flows (including IPv4 with tags) are not included in results.

For more information about the *Test Setup>SmartTracker* tab, refer to the online Help.

Summary of SmartTracker Capabilities and Limitations

Table 16-1 summarizes the capabilities and limitations of the SmartTracker test. For more information, see *"Limitations" on page 388*.

Table 16-1. SmartTracker Capabilities and LImitations Summary

If you want to test...	You can...	You cannot...
QoS (See *"How SmartTracker Tracks QoS" on page 390*.)	*For IPv4*, track the three ToS or DiffServ IP precedence bits.	Track more than 8,192 flows and ports if you use three bits for IPv4.
	For IPv6, track the traffic class bits.	Track more than 511 flows and ports if you use seven bits for IPv6.
VLAN (See *"VLAN Testing" on page 393*.)	Track the entire VLAN ID field (for up to 255 VLAN IDs) whose VLAN IDs range from 1 to 255 (least significant byte of the ID) and get flow-based results. Track the entire VLAN ID field (for up to 4095 VLANs IDs) whose VLAN IDs range from 1 to 4095 and get port-based results. You can still track the flow in results if you set up each flow with a unique VLAN ID. Test results show a separate tracked VLAN ID value listed for each port, which is equivalent to each flow. Track only the high-order range of bits in the VLAN ID and get flow-based results. This results in tracking the following specific high-numbered VLAN IDs (but none between these IDs): 256, 512, 768, 1792, 2048, 3840.	Track 4,095 VLAN IDs and get flow-based results. (You can get port-based results, but you will not be able to associate a specific flow with a VLAN ID unless a unique VLAN ID is assigned to each flow.) Track VLANs for POS cards.
MPLS	Track the three exp bits. However, since these bits do not coincide with a byte boundary, more than three bits are used, and this limits the number of flows in the test.	Track the entire MPLS label.
TCP or UDP	Track the entire source or destination TCP or UDP port number field with port-based results. Port number values can range up to 65,535. For flow-based results, track only a portion of the field.	Get flow-based results if the entire TCP or UDP source or destination field is tracked.

Test Methodology

The SmartTracker test uses the *Alternate Key* feature and new hashing algorithm of TeraMetrics and TeraMetrics-based modules. Instead of using the traditional SmartBits signature field in packets to track results, this feature uses an alternate field of up to 16 bits. Because the alternate key is used, cyclic flows or other user-specified fields are rendered trackable in results.

SmartTracker tracks results per flow and per port, or just per port if all 16 tracking bits are used. The more bits that are tracked in frames, the fewer flows can be tracked. If only port-based results are sufficient, up to 65,535 tracked values can be tracked.

Limitations

When setting up and running the SmartTracker test, keep in mind the following limitations:

Card limitations:

- The SmartTracker test is supported on all TeraMetrics and TeraMetrics-based modules except ATM.
- *For XLW-372xA, XFP-373xA, and POS-3518/3519 modules*:
  - Up to 15 (not 16) bits can be tracked. The 15th and most significant bit (MSB) of the tracking field is ignored for these cards. If the MSB needs to be tracked, you can set up the test to track the MSB in the right-most byte.
  - If mixed card types are used, mix only ports from each of these module types (XLW-372x, XFP-373x, and POS-3518/19). Tthese modules cannot be mixed with any other card types. This restriction does not apply to mixing other card types or to other tests.

Flow limitations:

- The test can track up to 32,768 SmartFlows. The maximum number of flows allowed for the test is reduced depending on the number of bits used for the alternate key field that the test is tracking. The larger the size of the field that is tracked, the fewer the the number of flows and/or variations in non-cyclic flows are allowed. To increase the number of flows, consider disabling some flows for the test. See *Figure 16-2 on page 389*.
- The test can track a maximum of 255 cyclic flows.
- When testing QoS on IPv6, the test can track a maximum of 511 SmartFlows.

Other limitations:

- The SmartTracker test can track a field up to two bytes long. If all 16 bits are used, it can track only ports (not flows) in results.
- The first two bytes of a packet are not trackable.
- The SmartTracker test does not support result sampling (on the *Options* tab).
- When the test is set up, bit selections must be contiguous.
- If predefined or custom filters are used in the SmartTracker test, the *Data Integrity* option (the default) in the SmrtFlow.ini file must be disabled to run the test.
- SmartTracker results do not include sequence tracking (in- and out-of-sequence packets).

Important: If more flows are set up than are allowed for the number of bits being used, you will see only receiving measurements for flow-based results. To see receiving as well as transmitting statistics for flows, reduce the number of bits being tracked or reduce the number of flows that are enabled.

For VLAN limitations and recommended test setup, refer to *"VLAN Testing" on page 393*.

Table 16-2 shows the number of SmartFlows allowed per number of bits used.

Note: The maximum number of trackable SmartFlows (non-cyclic) allowed applies to a single test setup, not to a single transmitting or receiving port.

Table 16-2. SmartTracker Maximum Flows Allowed per Bits Used to Track

Number of Bits Used	Maximum Trackable SmartFlows
1	32,768
2	16,384
3	8192
4	4096
5	2047
6	1023
7	511
8	255
9	127
10	63

Table 16-2. SmartTracker Maximum Flows Allowed per Bits Used to Track (continued)

Number of Bits Used	Maximum Trackable SmartFlows
11	31
12	15
13	7
14	3
15	1
16	0 (port-based tracking only)

How SmartTracker Tracks QoS

Quality of Service (QoS) provisioning categorizes and prioritizes traffic according to its importance of delivery. This ensures that priority traffic (for example, real-time applications such as streaming video) gets priority treatment during times of congestion.

Both ToS (*Figure 16-1*) and DiffServ (*Figure 16-2*) use the same precedence byte in the IP protocol header: the first three bits, 0 to 2.

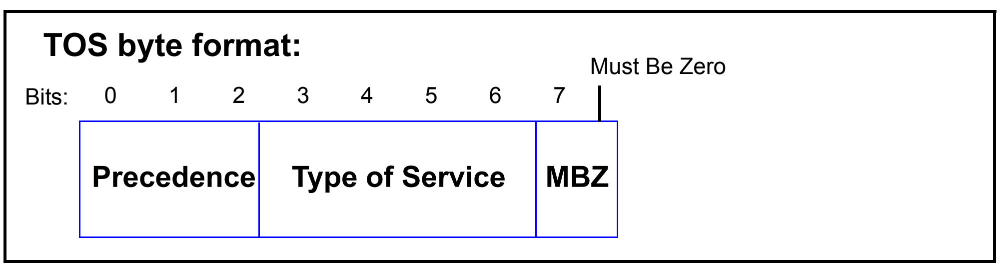

Figure 16-1. TOS Byte Format

Note: For IPv6, instead of tracking three bits in the ToS field, it tracks three bits in the traffic class field.

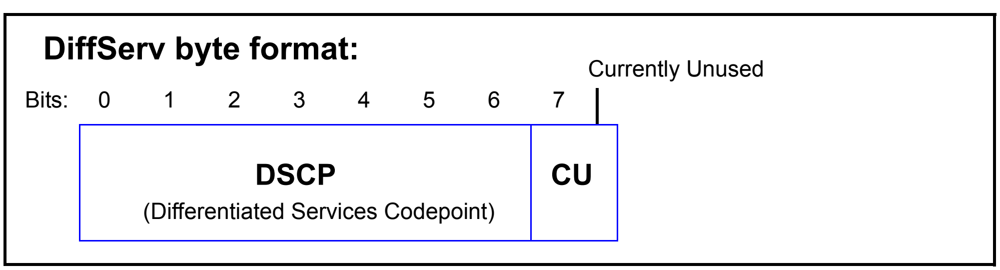

Figure 16-2. DiffServ Byte Format

A SmartTracker test set up to track QoS in IPv4 packets tracks the IP precedence bits in the ToS and DiffServ fields of the IP header. SmartFlow compares the values of these fields in the received frames to those originally transmitted. The original stream ID is compared with the ID in the received frame. This type of tracking is useful when testing a DUT that changes the IP precedence value based on certain packet characteristics.

For IPv6, the equivalent of the IP precedence bits is the traffic class octet.

IP Precedence Bits

The IP precedence field in the IP protocol header of a packet is a three-bit value that designates the relative priority of the packet. Values range from 0 to 7. Higher values are assigned to transactions that require expedited delivery.

Custom Tracking

When a SmartTracker test is set up to track a custom set of bits in test packets, the number of trackable SmartFlows is determined by the number of bits that are being tracked. See *Table 16-2, "SmartTracker Maximum Flows Allowed per Bits Used to Track," on page 389*.

For example, the three-bit experimental (exp) field in an MPLS label can be tracked. The field use is undefined, but in practice the field is used to implement a QoS mechanism within MPLS.

The TCP destination port or UDP destination port can also be tracked.

Tracking Cyclic Flows

To track cyclic flows, specify the offset for the source IP, destination IP, or MAC address in the frame. Each of these fields consist of four bytes.

It is possible to track the source IP address but vary the destination IP address (or vice versa) in cyclic flows.

Note: The test can track a maximum of 255 cyclic flows.

The results format varies according to how the test is set up. If you vary a field for a cyclic flow and track this same field, detailed results display a total line for each cyclic flow (tracked variation of the field). For example, *Figure 16-3* shows some of the Total lines for a custom field that was varied abnd tracked in the test. There is a total line for each variation of the field being tracked and for each iteration. The group line shows all variations in the group.

Name	Time	FrameSize	ILoad	Tx uniq Value	TxFrames	TxRate (%)	Rx uniq Value	RxFrames	RxR
		N/A	20.00000	244	5,341	0.39526	244	2,379	
		N/A	20.00000	245	5,341	0.39526	245	2,379	
		N/A	20.00000	246	5,341	0.39526	246	2,379	
		N/A	20.00000	247	5,341	0.39526	247	2,381	
		N/A	20.00000	248	5,341	0.39526	248	2,382	
		N/A	20.00000	249	5,341	0.39526	249	2,382	
		N/A	20.00000	250	5,341	0.39526	250	2,382	
		N/A	20.00000	251	5,341	0.39526	251	2,380	
		N/A	20.00000	252	5,341	0.39526	252	2,380	
		N/A	20.00000	253	5,341	0.39526	253	2,382	
		N/A	20.00000	254	5,341	0.39526	254	2,380	
		N/A	20.00000	255	5,341	0.39526	255	2,381	
A Group	12/23/03 15:53:02	N/A	N/A	3	5,341	0.39526	3	2,381	
		N/A	N/A	4	5,341	0.39526	4	2,382	
		N/A	N/A	5	5,341	0.39526	5	2,382	
		N/A	N/A	6	5,341	0.39526	6	2,381	
		N/A	N/A	7	5,341	0.39526	7	2,382	
		N/A	N/A	8	5,341	0.39526	8	2,381	
		N/A	N/A	9	5,341	0.39526	9	2,382	
		N/A	N/A	10	5,341	0.39526	10	2,383	
		N/A	N/A	11	5,341	0.39526	11	2,384	

Figure 16-3. SmartTracker Flow Detail Results with Destination IP Field Varied and Tracked

How Alternate Key Works

SmartFlow tracks flows by using the stream ID field, a proprietary field within each packet. The SmartTracker test may be run only on TeraMetrics and TeraMetrics-based modules, which use an alternate key to identify and track received flows. The alternate key consists of two bytes (16 bits) that start at a fixed offset from the beginning of the IP header. When a TeraMetrics-based module receives a packet, SmartFlow inserts either one or two bytes from the alternate key field into the stream ID, as shown in *Figure 16-4*. This becomes the place where histogram information is stored.

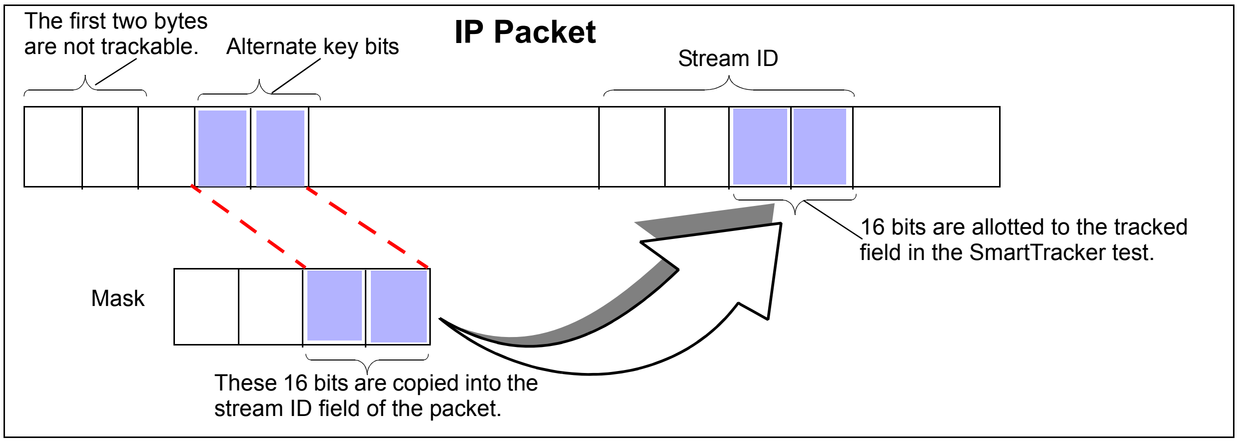

Figure 16-4. How Alternate Key Works

You can use a maximum of 16 bits to track a custom field. For example, if QoS is being tracked, three bits are for IP precedence and 13 bits are for the signature field.

How Cyclic Flows are Tracked

Instead of inserting the three IP precedence bits into the 16 bits of the alternate key, you can insert a portion of the source IP, destination IP, or MAC address into those bits. This enables you to track cyclic flows. Other custom fields can be tracked in a similar fashion.

VLAN Testing

You can perform limited VLAN testing—for VLAN stacking, VLAN leakage, or VLAN ID integrity—by using custom tracking of the VLAN ID. You can track up to 4095 unique VLAN IDs that range from 1 through 4095, with port-based results.

Plan the VLAN ID numbering according to what is required in test results:

- *If you want to track VLANs by flow as well as by port:*
  You can track up to 255 VLANS with VLANs IDs ranging from 1 to 255, which is the least significant (right) byte of the VLAN ID.

- *If you want port-based results that show VLANs of a receiving port:*
  You can track up to 4095 VLANs with VLAN IDs ranging from 1 to 4095, using both bytes. However, you will not be able to associate a specific flow with a VLAN.

Figure 16-5 shows the entire VLAN tag. The type field shows 8100, which indicates the presence of a VLAN tag. The VLAN ID portion itself is 12 bits long. Because SmartTracker resolves bits to the nearest byte, 16 bits (two bytes) is used. (This is the maximum number of bits allowed for the SmartTracker test.) Using the maximum number of bits limits the results to tracking VLAN IDs on a port-only basis, not by flow and port.

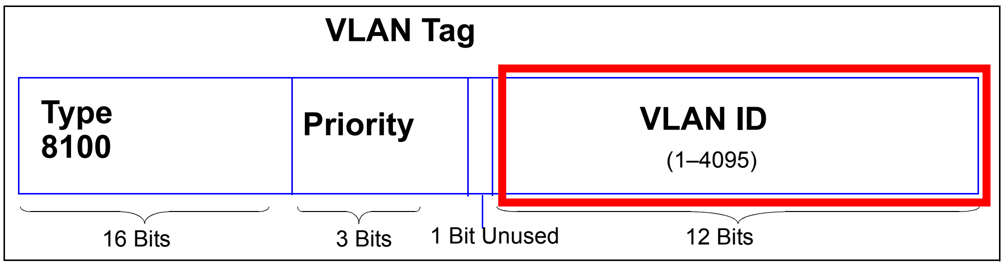

Figure 16-5. Entire VLAN Tag which Includes the VLAN ID

Tip: To verify that the VLAN IDs received from the DUT are the identical to the VLAN IDs sent in the test, set up the test to track the entire VLAN ID (the full range of VLAN IDs from 1 to 4095). This limits results to per-port, rather than per-flow, since tracking the entire VLAN ID field uses all 16 tracking bits. You can also do the following: Set up the flows so that each one has a unique VLAN ID. In test results, you will see a separate tracked VLAN ID value listed for each port (equivalent to each flow).

To track VLAN IDs on a per flow basis as well as by port, at the *Test Setup>SmartTracker* tab, specify fewer than 12 bits to track. Because of byte boundaries, tracking 12 bits actually uses all 16 bits. See *"Setup Tips for Flow and Port-based VLAN Tracking"* for more information.

Note: SmartTracker test results display tracked fields such as VLAN IDs or the Ethernet type field in decimal numbering. For example, if you track a VLAN ID of 25 (19 in hexadecimal), 25 is displayed in results.

Setup Tips for Flow and Port-based VLAN Tracking

Selecting the entire VLAN ID provides only port-based results, but by choosing to track a portion of the VLAN ID, you can track VLANs by both flow and port. When setting up the SmartTracker test, the portion of the VLAN tag to be tracked should determine the offset and bits that are selected. The more bits that are tracked, the more combinations of VLAN IDs can be tracked, but fewer flows can then be included in the test. (See *Table 16-2 on page 389*.)

Before setting up VLAN tracking, consider the following:

- Decide on the decimal numbers of the VLAN IDs to be tracked.
- Define the binary equivalent to these VLAN IDs.
- Decide how many IDs need to be tracked.

If the VLAN IDs to be tracked consist of two decimal digits (0–99), you can select the corresponding binary bits to track.

Important: To test VLAN IDs, if only a portion of the VLAN ID field is being tracked, the IDs on the DUT must be closely controlled. Make sure that you are only using the IDs that correspond to the bits set up on the *Test Setup>SmartTracker* tab. If other IDs are used, you can get unexpected VLAN IDs in results.

Tracking Lower-numbered VLAN IDs (1 to 255)

To track a range of lower-numbered VLAN IDs:

1. Click the **Test Setup>SmartTracker** tab.
2. Enter a **Offset (bytes)** value of **15**.

Figure 16-6 on page 396 shows an example of the bits to select at the *Test Setup>SmartTracker* tab to track VLAN IDs 1–255. VLAN IDs from 1 to 255 are decimal equivalents to the eight least significant bits of a VLAN ID. Tracking more bits of the VLAN ID uses more bits in the SmartTracker test setup, since the offset must be in bytes.

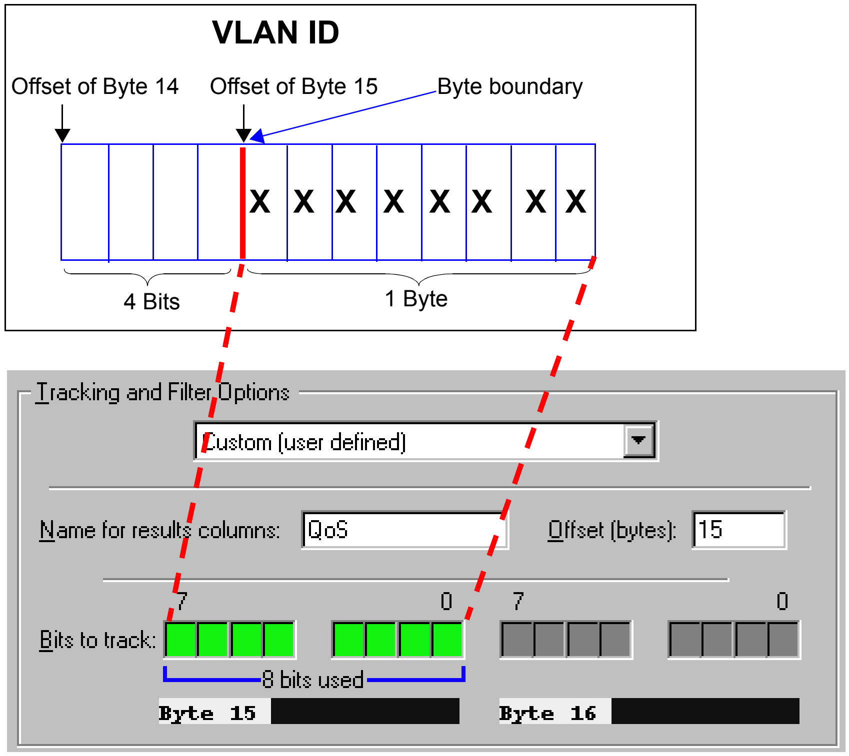

Figure 16-6. Example of VLAN ID Bits to Select in Order to Track VLAN IDs 1–255

Tracking Higher-numbered VLAN IDs (Greater than 255)

For higher-numbered VLAN IDs (greater than 255), you can still track flows in results by tracking only the high-order range of bits in the VLAN ID field. Tracking the high-order range results in tracking those *specific* high-numbered VLAN IDs, but not any IDs between them: 256, 512, 768, 1792, 2048, 3840. To track these bits, set the offset to 14.

Figure 16-7 illustrates the high-order bits to select that correspond to all those high-numbered VLAN IDs. Notice that when any of the four high-order bits are selected, the priority bits are also used (but not tracked) since the offset is in bytes and the priority bits are part of the same byte.

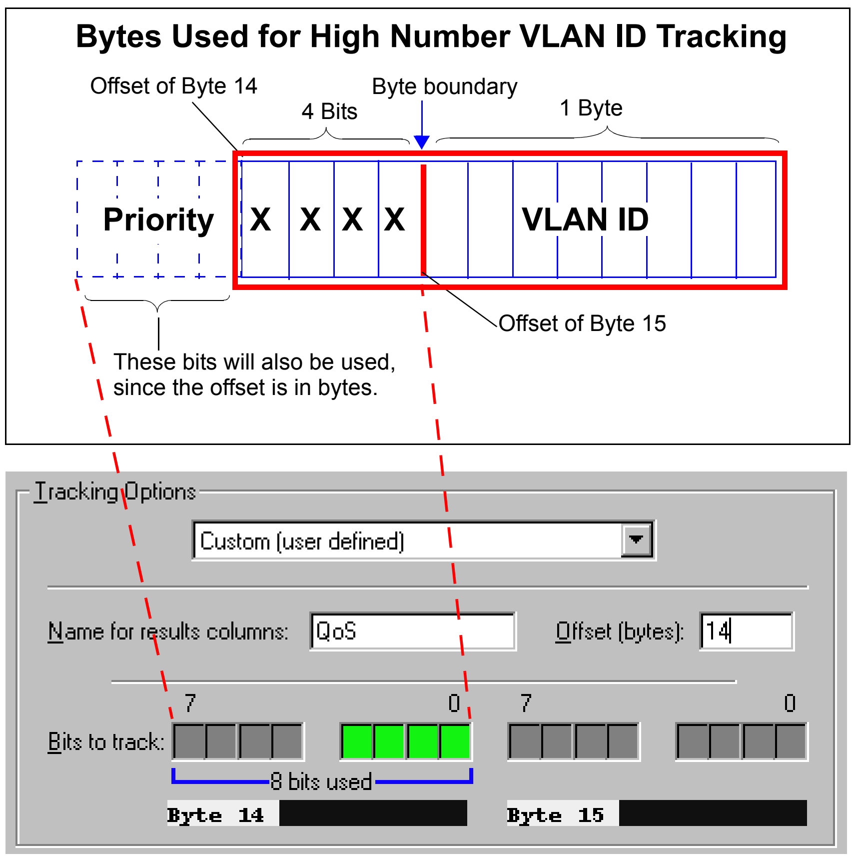

Figure 16-7. Tracking High-order VLAN ID Bits

The combination of bits that were set on the *Test Setup>SmartTracker* tab determine the combination of high order bits which track specific, known VLAN IDs.

The following diagrams show which bit(s) to set to obtain tracking of specific VLAN IDs.

To track VLAN ID 256, select the bit that corresponds to this high-order bit in the VLAN ID:

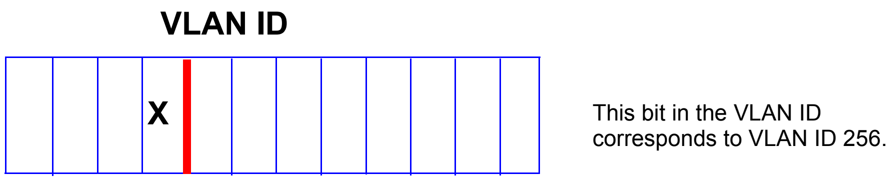

To track VLAN ID 512, select the bit setup that corresponds to this high-order bit in the VLAN ID:

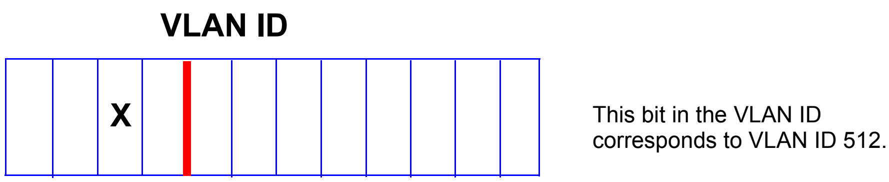

To track VLAN IDs 256, 512, and 768, select the bits that correspond to these high-order bits in the VLAN ID:

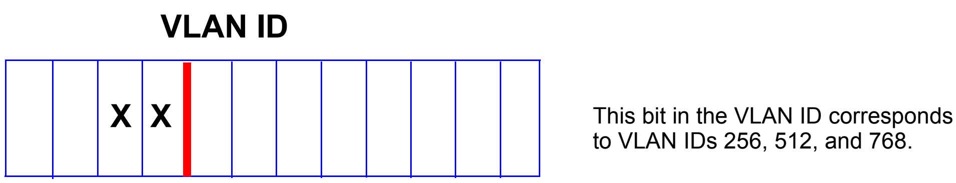

Testing VLAN Leakage

VLAN leakage can occur when the DUT:

- Places a VLAN tag on an untagged frame belonging to a different flow.
- Places the wrong VLAN tag on a tagged flow from another VLAN.

To test VLAN leakage, set up the test in the same manner as you would to track VLAN IDs. SmartFlow checks the proprietary signature field in each frame to ensure that it belongs with that VLAN ID.

Determining VLAN Leakage in Test Results

VLAN leakage has occurred if on the *Flow Detail* tab you see an *RxFrames* column with no flow name and no value in the *Tx Frames* column.

From a port-level perspective, VLAN leakage has occurred if on the *Port Detail* tab you see the VLAN ID received for a port on which it should not have been received. In this case, the VLAN tag leaked onto this port.

VLAN leakage can be partial or complete. If it is partial, then the receiving port has received some packets with that VLAN ID and some without it. If it is complete, then the port that was supposed to receive that VLAN ID has not received any packets with that VLAN ID.

Interpreting SmartTracker Test Results

Test results for the SmartTracker test appear in the main window as the test runs. Test results can be viewed in graphical or tabular form. The SmartTracker test provides the following types of results:

- Pie chart
- Rate chart (See *"SmartTracker Rate Chart Results" on page 401*.)
- Port detail (See *"SmartTracker Port Detail Results" on page 404*.)
- Flow detail (See *"SmartTracker Flow Detail Results" on page 405*.)
- Stray frames (See *"SmartTracker Stray Frames Results" on page 407*.)
- Port errors
- Packet rate.

Notes:

- All flows enabled for a test are displayed in results as transmitted frames, even if the filter excludes some traffic. For example, both IPv4 and IPv6 flows are present in a test but the IPv6 traffic is filtered out. The IPv6 flows still display in results as transmitted frames, but the value for the received frames of the excluded IPv6 traffic will be N/A.
- If no frames are received on the receiving port, the following result columns show N/A: RxFrames, RxFrames(%), Delay(Rx-Tx), Rx fps, Rx L3 bps, Rx L2 bps. A value of 0 in these columns indicates that no frames were received that match the tracked field. The value N/A indicates that tracking is not applicable (i.e., there were no frames to track).
- When no frames are received on any of the receive test ports, SmartFlow does not show the receive-side result columns at all.
- Results format can vary depending on whether or not you used frame size automation during the test. See *"Using Frame Size Automation (Global for All Flows)" on page 236*.

Tip: If test results differ from what is expected, run a sample iteration, and capture and analyze the received packets to see if the DUT altered the packet structure.

SmartTracker Pie Chart Results

SmartTracker pie chart results display the ratio of frames received for each value of the tracked field, relative to the other tracked field values in an iteration. For example, if you are tracking three levels of QoS with IPv6, the pie chart shows three sections, one for each QoS value. (See *Figure 16-8*.) Each segment of the pie chart represents a total for a tracked value and corresponds to each total line for the *Rx Tracked Value column* on the *Flow Detail* tab.

The maximum number of tracked entities or values that can be displayed is 20. If there are more than 20 entities in the test, the first 20 entities are displayed.

If frame size automation was used, select the desired frame size at the top of the *Results* tab.

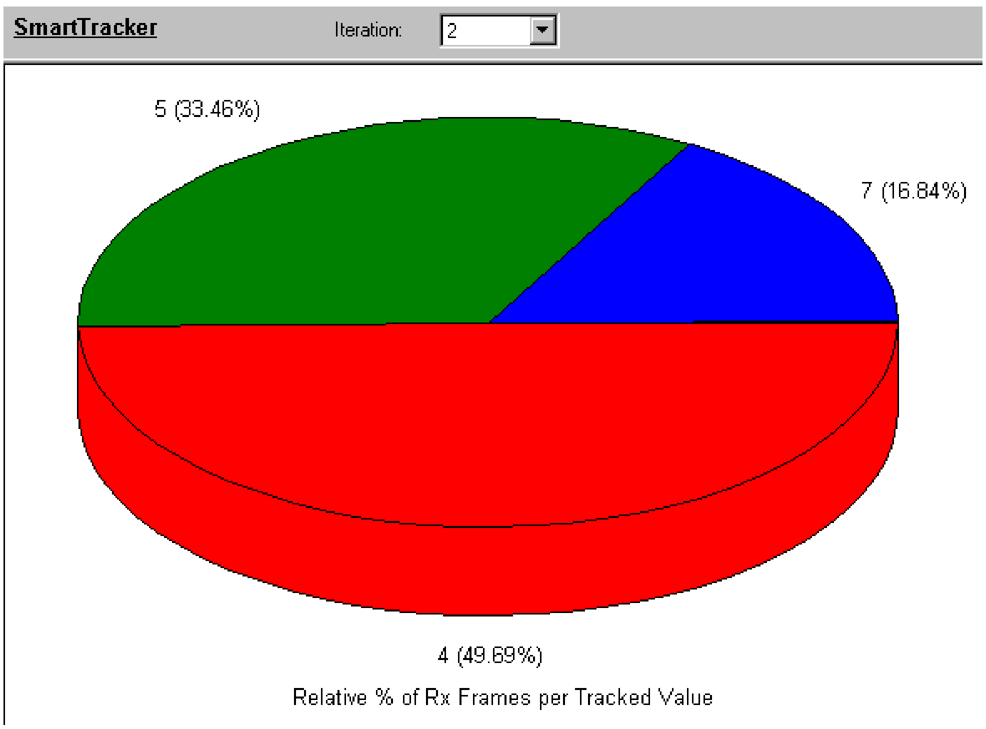

Figure 16-8. SmartTracker Pie Chart Results

SmartTracker Rate Chart Results

Rate chart results display (for each group in an iteration) the number of frames received for each value of the tracked field. For example, if three levels of QoS are tracked, the rate chart shows the receive rate for three sections—one for each QoS value.

For a description of each field in SmartTracker test rate chart results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Selecting the Unit for Display

Some DUTs have interfaces with bandwidth that is less than the rate at which the SmartBits card/module is transmitting or the interface only supports a specific bandwidth. To compare the receive rate of the DUT to the rate being used in the test, specify the units by which to view the rate chart results. To select a unit, right-click anywhere on the chart. You can view rate chart results in any of these units:

- Megabits per second
- Kilobits per second
- Bits per second
- Frames per second
- Tracked value (%).

The rate units refer to the receiving rate and are presented in scientific notation if the value goes beyond six digits. (See *Figure 16-9* for an example.) For information about scientific notation, see *"Scientific Notation in Chart Results" on page 81*.

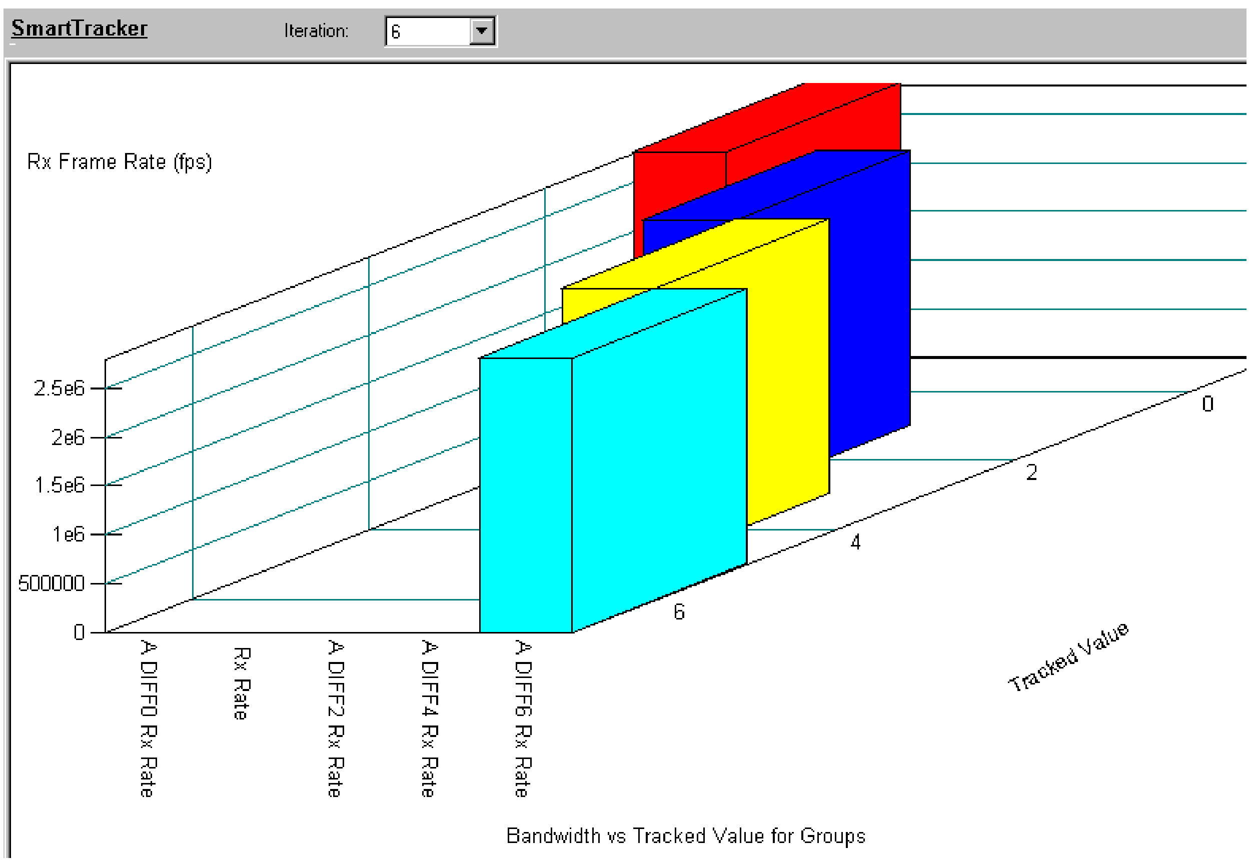

Figure 16-9. SmartTracker Rate Chart Results

If the Rate Chart is Empty

The rate chart depends upon flow statistics to present result histograms. If the test setup creates more flows than SmartTracker is able to track, this error message displays:

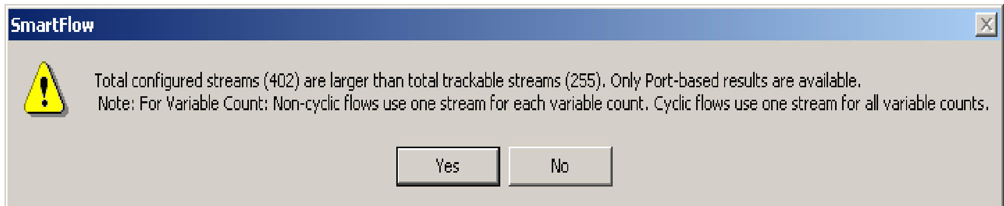

In this example, SmartTracker is able to track 255 streams, fewer than the 402 that have been configured. It provides port-based results, but will not fill the rate chart.

In response, reduce the number of configured streams to match the displayed "total trackable streams" in the message. Then run the test again.

SmartTracker Port Detail Results

The *Port Detail* tab displays detailed information for a tracked field based on only the receiving port, as shown in *Figure 16-10*.

For information about how to determine if VLAN leakage occurred, see *"Determining VLAN Leakage in Test Results" on page 399*.

Note: When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

For a description of each field in the SmartTracker test port detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

SmartTracker   Page 1 of 1   Next

Name	Time	QoS Value (Dec)	TxFrames	RxFrames	TxRate (%)	RxRate (%)	Tx fps	Rx fp
Totals (All Ports)	10/08/04 13:28:28	0	422,297	422,297	50.00000	50.00000	42,229	[illegible]
		2	422,297	422,297	50.00000	50.00000	42,229	[illegible]
Subtotal	10/08/04 13:28:28	N/A	844,594	844,594	N/A	N/A	84,458	[illegible]
SMB 2 1-2	10/08/04 13:28:28	0	0	422,297	0.00000	50.00000	0	[illegible]
		2	0	422,297	0.00000	50.00000	0	[illegible]
Subtotal	10/08/04 13:28:28	N/A	0	844,594	N/A	N/A	0	[illegible]
SMB 2 1-1	10/08/04 13:28:28	0	422,297	0	50.00000	0.00000	42,229	
		2	422,297	0	50.00000	0.00000	42,229	
Subtotal	10/08/04 13:28:28	N/A	844,594	0	N/A	N/A	84,458	
Totals (All Ports)	10/08/04 13:28:48	0	844,594	844,594	50.00000	50.00000	84,459	[illegible]
		2	844,594	844,594	50.00000	50.00000	84,459	[illegible]
Subtotal	10/08/04 13:28:48	N/A	1,689,188	1,689,188	N/A	N/A	168,918	[illegible]
SMB 2 1-2	10/08/04 13:28:48	0	0	844,594	0.00000	50.00000	0	[illegible]
		2	0	844,594	0.00000	50.00000	0	[illegible]
Subtotal	10/08/04 13:28:48	N/A	0	1,689,188	N/A	N/A	0	[illegible]
SMB 2 1-1	10/08/04 13:28:48	0	844,594	0	50.00000	0.00000	84,459	
		2	844,594	0	50.00000	0.00000	84,459	
Subtotal	10/08/04 13:28:48	N/A	1,689,188	0	N/A	N/A	168,918	
Totals (All Ports)	10/08/04 13:29:09	0	1,266,891	1,266,891	50.00000	50.00000	126,689	[illegible]

Figure 16-10. SmartTracker Port Detail Results

SmartTracker Flow Detail Results

The *Flow Detail* tab (*Figure 16-11*) displays results by flow and/or group. It displays detailed information (per iteration) about each value of the tracked field, by flow and/or group.

Notes:

- This tab does not display received statistics for flows if the tracked field is larger than 15 bits. In this case, only statistics for the transmitting port of the flow are displayed, but detailed port-based results are available on the *Port Detail* tab.
- When reporting the %load, Smartflow includes both IFG and preamble in the calculation.This generates a load value relative to the full transmission rate and not relative to the maximum data rate.

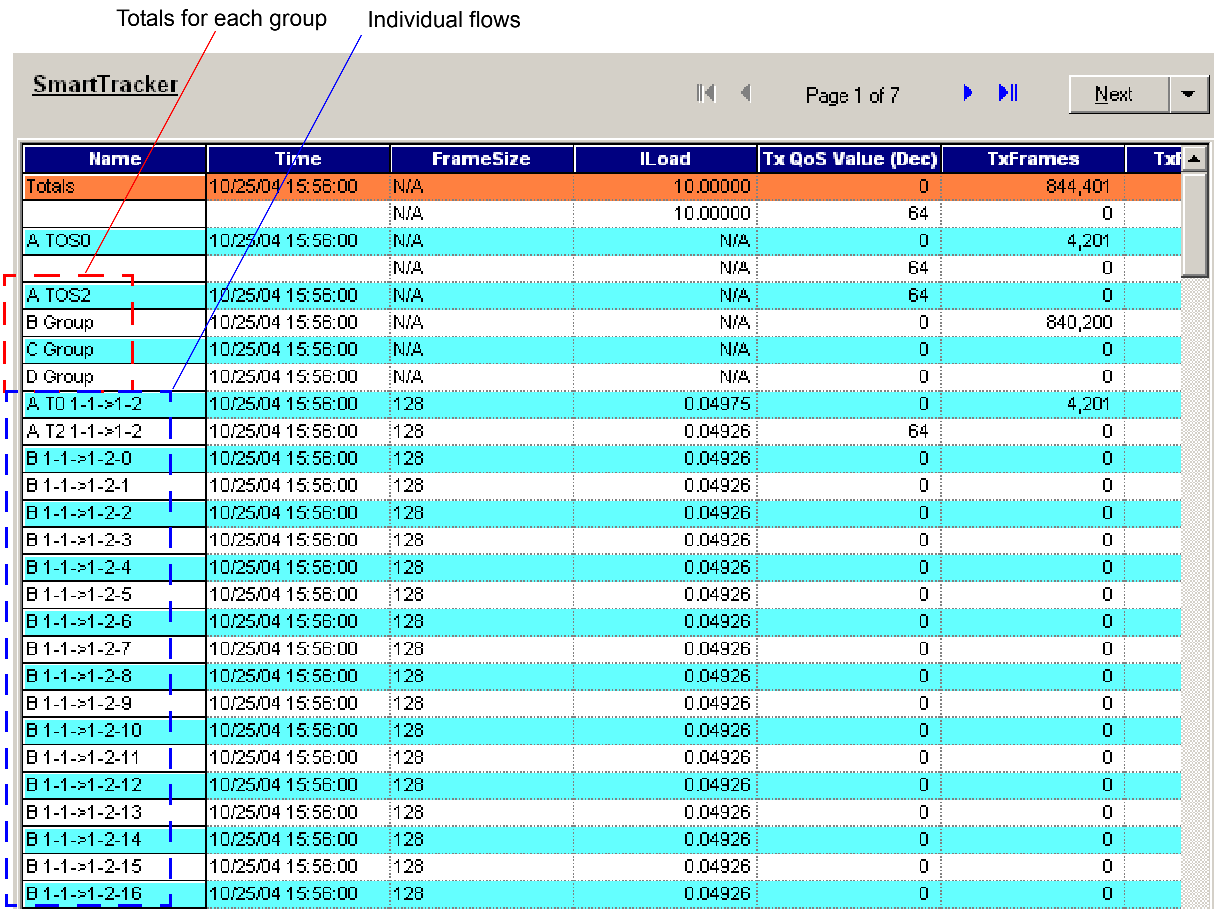

Name	Time	FrameSize	ILoad	Tx QoS Value (Dec)	TxFrames	TxF
Totals	10/25/04 15:56:00	N/A	10.00000	0	844,401	
		N/A	10.00000	64	0	
A TOS0	10/25/04 15:56:00	N/A	N/A	0	4,201	
		N/A	N/A	64	0	
A TOS2	10/25/04 15:56:00	N/A	N/A	64	0	
B Group	10/25/04 15:56:00	N/A	N/A	0	840,200	
C Group	10/25/04 15:56:00	N/A	N/A	0	0	
D Group	10/25/04 15:56:00	N/A	N/A	0	0	
A T0 1-1->1-2	10/25/04 15:56:00	128	0.04975	0	4,201	
A T2 1-1->1-2	10/25/04 15:56:00	128	0.04926	64	0	
B 1-1->1-2-0	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-1	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-2	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-3	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-4	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-5	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-6	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-7	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-8	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-9	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-10	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-11	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-12	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-13	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-14	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-15	10/25/04 15:56:00	128	0.04926	0	0	
B 1-1->1-2-16	10/25/04 15:56:00	128	0.04926	0	0	

Figure 16-11. SmartTracker Flow Detail Tab

For information about how cyclic flows appear on this tab and sample results, see *"How Cyclic Flows are Tracked" on page 393*.

For information about how to determine if VLAN leakage occurred, see *"Determining VLAN Leakage in Test Results" on page 399*.

For a description of each field in the SmartTracker flow detail results as well as the customized detailed reporting function, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Flow Detail Example of a Delta Between Tx and Rx Values

The *Delta* column on the *Flow Detail* tab shows the number of received frames whose tracked field value varies from the value that was transmitted for the flow.

For example, a test is set up to track the VLAN ID in frames. Assume that three frames are transmitted in flow A, each with a VLAN ID of 10. However, only one of the three frames arrives with a VLAN ID of 10. *Figure 16-12* illustrates this situation in more detail.

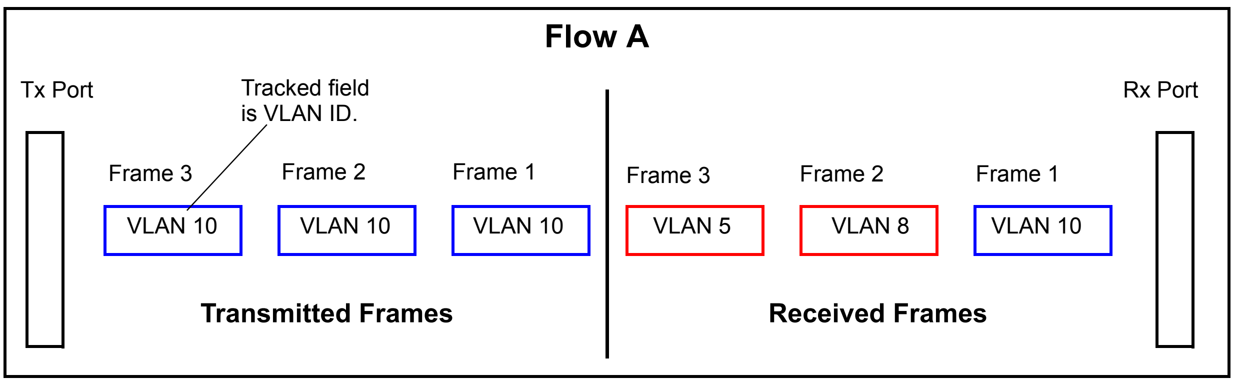

Figure 16-12. Example of Delta between Transmitted and Received Tracked Field Value

Based on the example in *Figure 16-12*, the related fields on the *Flow Detail* tab would look like this for flow A:

Tx VLAN Value	TxFrames	Rx VLAN Value	RxFrames	Delta
10	3	10	1	– 2
8	0	8	1	+1
5	0	5	1	+1

SmartTracker Stray Frames Results

The SmartTracker stray frames results table shows (per port) frames in the snapshot that were received but were not destined (expected) for that port. (See *Figure 16-13*.) It also lists each flow (not SmartFlow) in the snapshot containing the stray frames received by that port. Stray frames are identified in two ways:

- They contain the SmartBits signature field. This indicates that they are SmartBits test frames, not frames generated by the DUT.
- The destination IP address of the flow does not match any of the IP addresses associated with the port that received the frames.

Stray frames results list only those flows of a SmartFlow whose frames were misdirected and only if the *Record stray frames* checkbox on the *Options>General* tab is selected. If a port did not receive any stray frames from a flow, that flow does not appear on the *Stray Frames* tab, even if other flows in the same SmartFlow are listed.

Each row in the *Stray Frames* tab represents one flow of a SmartFlow. The source IP address, destination IP address, and SmartFlow name uniquely identify the flow.

Number of Records in Stray Frames Results

When *Record stray frames* checkbox is selected on the *Options>General* tab, SmartFlow collects records for up to 50 stray streams (frames) on each port. You can adjust this default limit by editing the SmrtFlow.ini file, which is located in the SmartFlow program directory. Refer to the online Help for instructions to edit this file.

For a description of each field in SmartTracker test stray frames results, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

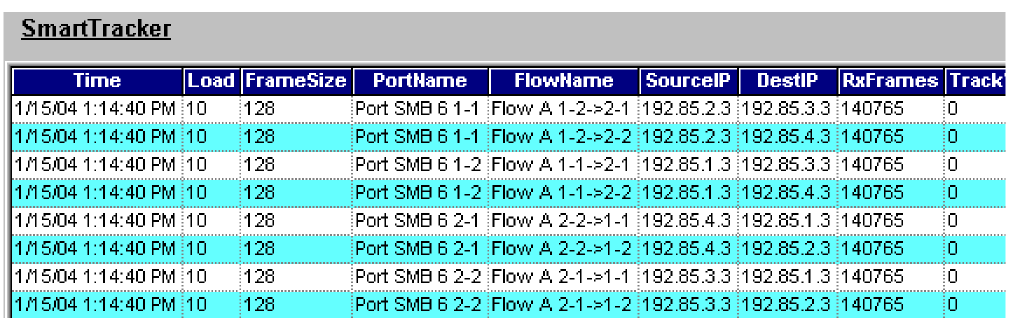

SmartTracker

Time	Load	FrameSize	PortName	FlowName	SourceIP	DestIP	RxFrames	Track
1/15/04 1:14:40 PM	10	128	Port SMB 6 1-1	Flow A 1-2->2-1	192.85.2.3	192.85.3.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 1-1	Flow A 1-2->2-2	192.85.2.3	192.85.4.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 1-2	Flow A 1-1->2-1	192.85.1.3	192.85.3.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 1-2	Flow A 1-1->2-2	192.85.1.3	192.85.4.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 2-1	Flow A 2-2->1-1	192.85.4.3	192.85.1.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 2-1	Flow A 2-2->1-2	192.85.4.3	192.85.2.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 2-2	Flow A 2-1->1-1	192.85.3.3	192.85.1.3	140765	0
1/15/04 1:14:40 PM	10	128	Port SMB 6 2-2	Flow A 2-1->1-2	192.85.3.3	192.85.2.3	140765	0

Figure 16-13. SmartTracker Stray Frames Results

Chapter 17

MPLS Testing

This chapter contains information about the MPLS features of SmartFlow and general information about MPLS technology.

In this chapter...

What is MPLS?

Multiprotocol Label Switching (MPLS) is a standards-based network management solution for large scale IP networks. Originally presented as a means of speeding router forwarding performance, MPLS traffic engineering capabilities are now considered its greatest strength.

The essence of MPLS is the Label Switched Path (LSP) to which specified packets are assigned using a special 32-bit label header. The LSP is a logical, predetermined tunnel (path) through a specified medium. Labeling is a shorthand representation of the IP header that permits network managers to dictate the path that traffic takes through a network. Labels can distinguish routing information as well as application type or service class.

MPLS works with IP, ATM, and frame relay network protocols. It allows most packets to be forwarded at the Layer 2 level. Label switching replaces conventional longest-address-match forwarding with a more efficient label swapping algorithm.

MPLS Label Stack Entry

MPLS packets are identified by the presence of a label stack between the layer 2 header and the layer 3 header in the packet. A label stack is composed of one or more label stack entries.

A label stack entry is diagrammed in *Figure 17-1* below. A label stack entry is a 32-bit value made up of four fields. The 20-bit label field contains a value that is used to identify a particular LSP. Labels have local significance only and are mapped in a one-to-one relationship with established LSPs on the local port.

Figure 17-1. MPLS Label Stack Entry

Label values less than 16 are reserved for special functions. The 3-bit experimental (exp) field has an undefined use by definition, but is actually used to implement a QoS mechanism within MPLS. The 1-bit, S-bit indicates whether or not the current label stack entry is the final entry in the label stack. Only one label stack entry within a label stack has this bit set to one, and this label stack entry always is the last label stack entry before the IP header. Finally, the IP Time-to-Live (TTL) information is also maintained in the label stack entry. (See *Figure 17-1*.)

How the TTL Byte of the Stack is Used

The IP TTL is duplicated in each MPLS label stack because information in the IP header of a classified packet is not modified as the packet moves through the LSP. If LSRs were to operate on the TTL in the IP header, each LSR would need to locate this information, change the TTL, and update the checksum. When a packet is placed into an LSP by the ingress LER, the IP TTL information is copied into the label stack entry. As the packet moves through the LSP, each LSR adjusts the TTL information in the label stack entry. When the packet reaches the egress LER, the TTL information is copied back into the IP header and the IP header checksum is recomputed one time.

It is possible to have multiple label stack entries within the label stack. If a packet has multiple label stack entries, the label nearest the IP header is the trunking LSP. An LSR examines only that label nearest the IP header. When the packet reaches the termination of the trunk LSP, this label is removed and other labels remain. The TTL information is copied from the trunk to the flow LSP, and the packet is resubmitted to the MPLS engine based on the new flow LSP.

Label Switching Routers

An MPLS-enabled device, called a Label Switching Router (LSR), supports standard IP control (routing protocols, RSVP) plus a label swapping component. There are three types of LSRs in an MPLS network:

- *Ingress LERs* (also called edge LSRs or ingress LSRs)
  Receive IP packets from an IP network and send out MPLS packets. Responsible for initial packet processing, classification, and label application.
- *Transit LSRs*
  Receive MPLS packets and forward MPLS packets in an MPLS network.
- *Egress LERs* (also called edge LSRs or egress LSRs)
  Receive MPLS packets and forward IP packets to an IP network.

Figure 17-2 shows the three types of LSRs in a typical MPLS network configuration.

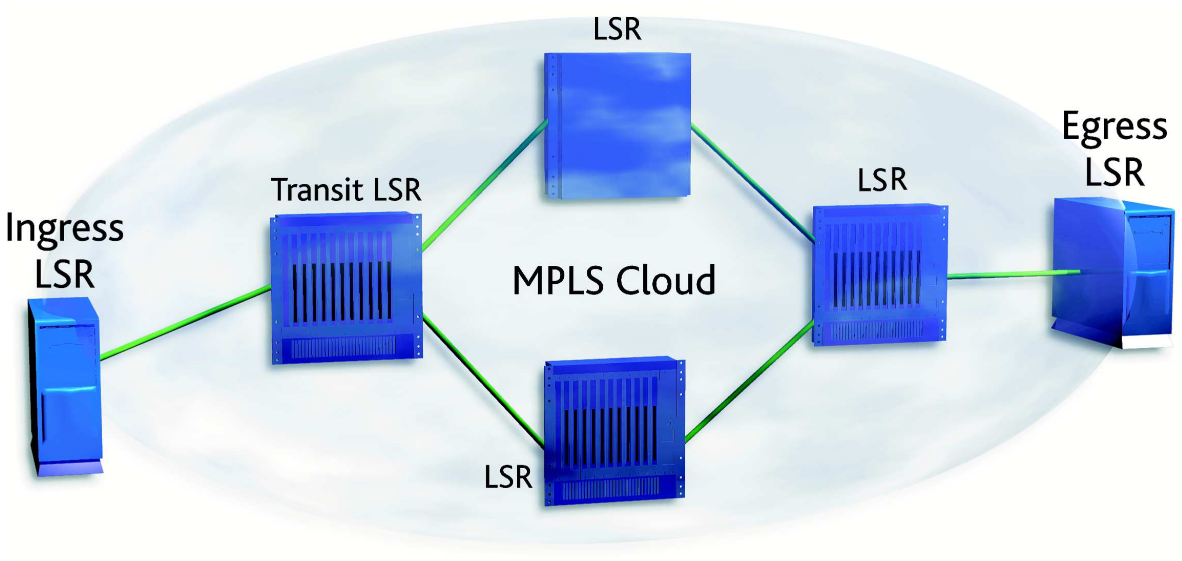

Figure 17-2. Typical MPLS Network

Ingress LER

An ingress LER converts IP packets into MPLS packets (by inserting a shim, or MPLS label stack) and forwards the packets toward their destination. Ingress LERs are the starting point of all LSPs, and therefore initiate all of the label distribution protocol signaling.

The ingress LER examines each IP packet and classifies the packet as a member of a specific LSP or route. Because it is desirable to classify packets based upon multiple criteria such as destination IP address, source IP address, and DiffServ code point, the ingress LER usually dedicates much more processing to packet classification than a traditional router. Once it classifies a packet, it applies the MPLS label to the packet and forwards the packet out the appropriate LSP. Transit LSRs use the label for the forwarding decisions.

Ingress LER functionality is illustrated in *Figure 17-3 on page 413*.

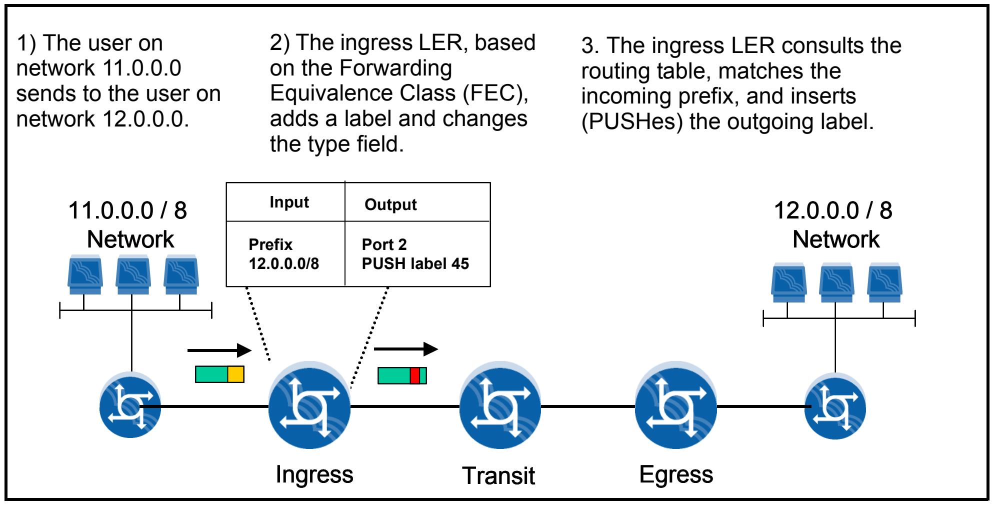

Figure 17-3. Ingress LER Functions

Note: Not every packet that enters an ingress router is forwarded as an MPLS packet.

Because MPLS uses traditional IP routing as a base, an MPLS network not only forwards labeled packets on LSPs, but routes non-labeled packets using standard routing rules as well. Only a fraction of the traffic in an MPLS network actually has a label; the remaining traffic is routed through the LSRs using normal IP routing techniques. This point is critical to understanding MPLS performance because traditional routing performance is a key component to overall MPLS network performance.

Transit LSR

When a transit LSR receives a labeled packet, it reads the label, consults its forwarding table, then replaces the label, and retransmits the packet as a labeled packet.

Transit LSR functionality is illustrated in *Figure 17-4 on page 414*.

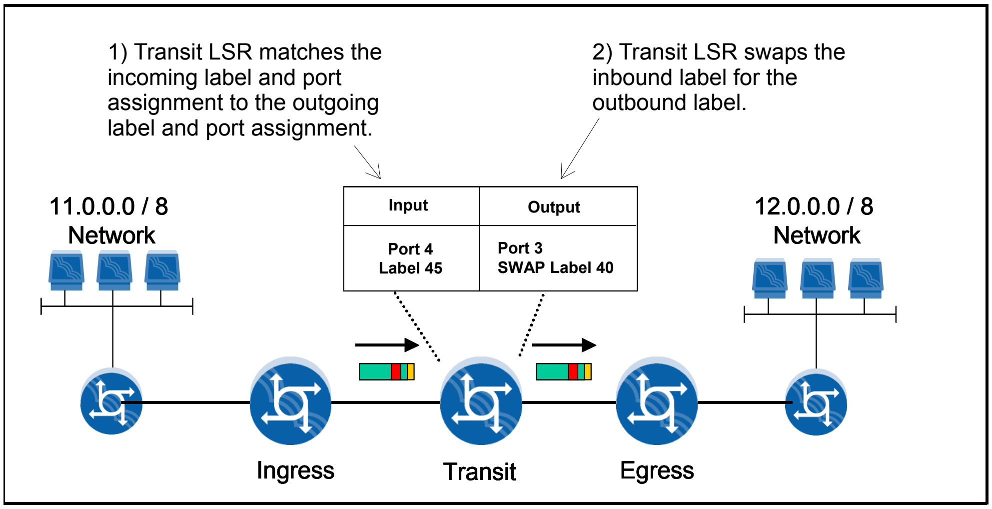

Figure 17-4. Transit LSR Functions

Because both IP and MPLS-based traffic are active in an MPLS network, the transit LSR also routes IP packets. Transit LSRs usually have the highest performance of all of the LSRs in the MPLS network.

Egress LER

An egress LER receives labeled packets, removes the MPLS label, copies TTL information from the MPLS label into the IP header, recomputes the IP checksum, and then routes the packet according to its IP header information.

Egress LER functionality is illustrated in *Figure 17-5 on page 415*.

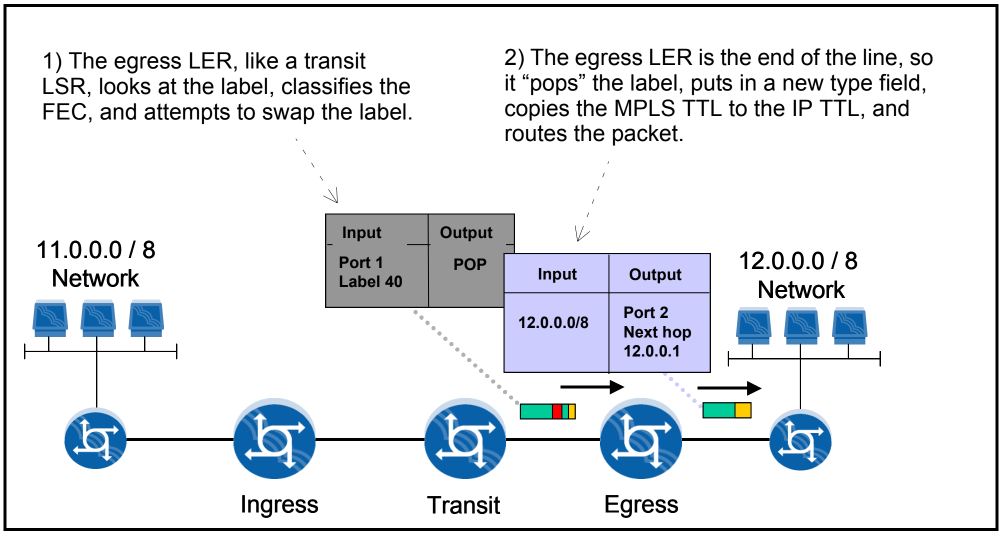

Figure 17-5. Egress LER Functions

Labeling, a shorthand representation of the IP header, permits network administrators to dictate the path that traffic takes through a network. Labels can distinguish routing information as well as application type or service class.

MPLS works with IP, ATM, and frame relay network protocols. It allows most packets to be forwarded at the Layer 2 level. Label switching replaces conventional longest-address-match forwarding with a more efficient label swapping algorithm.

Label Distribution Protocols

Label Distribution Protocol (LDP) is the general category of signaling protocols that can be used in MPLS to communicate labels and their meanings to MPLS-enabled devices called Label Switching Routers (LSRs). An LSR uses LDP to communicate its intention to use a particular label value on a specific interface. Working in conjunction with interior gateway protocols (BGP, OSPF, IS-IS), LDP establishes LSPs by assigning labels in edge and transit devices. As network routers participate in interior gateway protocols, routing tables are built.

Note: LDP is also the name of a specific protocol that was the first type of LDP. This is explained in more detail later in this documentation.

The QoS and path to be associated with an LSP are communicated at the time the LSP is created. Once an LSP is successfully created, the ingress LER can bind a route, or multiple routes, to an LSP. Only the ingress LER knows which packets will be placed in a particular LSP.

Manually Configuring LSPs

LSPs can be configured manually as a static route, without using an MPLS signaling protocol. It requires a manual configuration at the DUT at each hop along the path. A network administrator must log into each LSR to be traversed in the creation of the LSP. On each LSR, the administrator must specify the quality of service to be used with the LSP, the input port and label value, and the output port and label value. Manual configuration can be useful when working around a temporary limitation in the network, but it is impractical to use when scaling the network to thousands of LSPs.

Many LSR manufacturers do not support manual LSP configuration. LDPs allow the service provider to scale to many more LSPs.

Types of LDPs

Four different protocols currently exist for label distribution:

- Label Distribution Protocol (LDP) – A simple LDP that allows LSRs to communicate the labels to be used when forwarding traffic. The specification of any QoS information is not provided.
- Constraint-based Routing using LDP (CR-LDP) – Extends LDP so that it can be used to signal QoS information as well as label information.
- Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE) – Alternately used to distribute QoS and label-binding information. It allows you to specify attributes such as bandwidth and minimum/maximum frame size. This is the protocol that SmartFlow currently supports.
- Border Gateway Protocol Version 4 (BGP-4) – Enhanced to support use of label-bindings without any QoS information.

For a comparison of LDPs, see *“Comparison of LDPs” on page 417*.

Using LDP to Establish LSPs

Although LSPs can be configured manually, MPLS is easier to scale with LDP. LSP information is specified only once, reducing the chance for error and increasing the number of LSPs that can be established in a short time.

To use an LDP, an administrator typically logs into the management console on the ingress LER and specifies the desired QoS and path. The ingress LER initiates signaling toward the destination to set up the requested LSP. Once established, the administrator is notified and the LSP is available for forwarding traffic.

The network administrator can set up the network to monitor itself and create LSPs dynamically based upon current network conditions. While this approach has scalability advantages, it also presents the risk of creating a storm of signaling messages when the network is unstable. In this case, the administrator creates a policy on the ingress LERs in the network and allows the ingress LERs to automatically initiate signaling and to automatically bind routes to LSPs.

LDP Performance

The performance of the LDPs becomes a critical issue when trying to determine the overall performance of an MPLS network. While large numbers of LSPs will not usually be established or torn down as part of normal MPLS operation, these events can happen during periods of instability in the underlying routing protocols.

In addition, the amount of signaling required to indicate that an established LSP is still active can also limit the total number of LSPs that an LSR can support. While it is rare for an LSR to completely stop functioning, when an LSR is near its capacity, its LSPs will begin to fail and then be established automatically some time later. The effect is very similar to route flapping in traditional IP routing.

Comparison of LDPs

Because no single LDP satisfies all applications (voice, video, or data), a variety of LDPs target specific applications.

LDP:
The original Label Distribution Protocol (LDP*)* is specified by the IETF. It is used to distribute label bindings between adjacent LSRs, but no capability exists to specify any QoS-related issues such as reservable bandwidth or a specific path. The LSPs are created in a hop-by-hop manner and follow the paths that packets take when following the routing protocols. LDP tends to be used in conjunction with BGP-4 to allow service providers to increase the scale of their VPN solutions.

BGP-4:
BGP-4 can be extended to associate a label binding with a particular route. Both BGP-4 and LDP use TCP as a transport protocol. Because TCP is a reliable transport, these protocols are not concerned with retransmission or packet integrity.

LDP and BGP-4:
For the VPN solution, LDP is used to distribute labels for the LSP trunks between all of the edge routers in the service provider network. BGP is used to identify the individual VPN flows within the trunks so that one trunk can carry traffic from multiple customers without mixing the traffic. The core of the network only needs to be aware of the LSP trunks and can manage traffic at the trunk level. A provider's edge router will only be aware of the LSPs that terminate within the LSR. This approach should work well to scale VPN solutions, but it does not address customers who want a specific QoS level.

CR-LDP or RSVP-TE:
CR-LDP or RSVP-TE are used to satisfy QoS concerns. CR-LDP extends LDP by adding the capability to specify a path to the destination and to request some QoS parameters such as reservable bandwidth. RSVP-TE extends the RSVP protocol, which was designed for requesting QoS flows, to add label distribution. These two protocols perform the same task in an MPLS network and therefore a service provider rarely enables both in its network.

CR-LDP and RSVP-TE use different transport protocols. CR-LDP uses TCP as its transport, while RSVP-TE uses IP or UDP. Because IP and UDP are unreliable transports, RSVP-TE must concern itself with reliability issues. Retransmit timers are built into the protocol and refresh messages must be periodically sent to indicate that the LSPs are still established.

QoS and MPLS

The basis of Quality of Service (QoS) is the principle that transmission rates, error rates, and other network characteristics can be measured, improved, and (to some extent) guaranteed in advance. QoS is of particular concern for the continuous transmission of high-bandwidth multimedia information. Transmitting this kind of content dependably is difficult in public networks using ordinary "best effort" protocols.

IP routing has QoS limitations. If an ISP would like to offer one customer better service than is offered to another customer, routers must perform a longest prefix match not only on the destination IP address, but on the source IP address as well. This process doubles the time required for a standard IP route. Furthermore, each router in the forwarding path must pay the double lookup penalty. The requirement to perform two lookups also creates difficulties in using transit networks where the user does not have administrative control.

In addition, IP provides no standard mechanism to pass QoS information to all routers in the network. As a result, the network administrator must configure QoS for the user on each router.

The MPLS suite of traffic engineering and QoS components, bundled with a signaling mechanism like RSVP, provides a scalable solution because of these features:

- As a part of LDP, Forwarding Equivalence Classes (FECs) can be used to specify the LSP assigned to selected data. FECs permit LSRs to provide basic QoS to an LSP and the data it carries.
- Explicitly routed LSPs allow traffic to be engineered to travel through the network, avoiding congestion (according to QoS requirement, security classification, etc.) This capability within MPLS helps the network manager engineer a network that sustains QoS. LSPs can be manually created through the network to ensure QoS guarantees.
- Label switching used with QoS attributes allows different classes of ISP access service to be defined. QoS controls are easily communicated using the MPLS packet. MPLS devices allow QoS to be allocated to individual IP flows, then mapped to appropriate classes of service.

How MPLS Improves QoS

MPLS QoS solutions rely heavily on the ingress LERs and the traffic engineering capabilities of MPLS to achieve the desired results. The ingress LERs are responsible for detecting traffic from the preferred customer by examining the source IP address. The packet is then examined to see if it is a video packet or if it is data traffic. Finally, the packet's destination IP address is examined. Based upon this information, the packet can be placed into one of the LSPs listed in *Table 17-1*.

Table 17-1. LSP Determination Based on QoS

If the packet is ...	Then it is placed into ...
Video traffic	A guaranteed bandwidth pipe reserved exclusively for that customer's video traffic. This LSP can be created using CR-LDP or RSVP-TE to specify the necessary QoS to the transit and egress routers.
From the preferred customer	One of the LSPs reserved for that customer between the provider's edge routers. The packet is placed into the LSP whose egress point is closest to the packet's ultimate destination.
Video traffic from a regular customer	A bandwidth engineered LSP that is shared between multiple customers.

Using MPLS to roll out QoS in a network has a very significant benefit. The transit and egress LERs are not required to examine the packet in depth to understand what QoS levels to apply. Because the information is signaled when the LSPs are created, the LSRs can set up the appropriate queuing resources when the LSP is established and simply submit the labeled packets to their appropriate queues.

This same feature comes into play when one considers a service provider's desire to roll out a new QoS scheme in the existing MPLS network. In this case, the service provider needs to simply reconfigure the ingress LERs to create the appropriate LSPs and to identify the packets. The transit LSRs in the network are not aware of the fact that the policy has changed. With this approach to QoS, it is possible to roll out QoS incrementally without changing all of the firmware in all of the routers at one time, as is required in IP routing.

What is the MPLS Feature in SmartFlow?

The MPLS feature in SmartFlow works with TeraMetrics and TeraMetrics-based cards. It allows you to test the ability of a DUT to recognize MPLS labels and route MPLS traffic. It is also used for testing QoS with MPLS. The MPLS feature utilizes all of the existing features of SmartFlow plus the following:

- Supports up to 5,000 LSPs per test.
  Supports up to 511 LSPs per transmitting port, and up to 2047 for TeraMetrics and TeraMetrics-based LAN-332x ports.
- Supports the RSVP-TE signaling protocol.
  TeraMetrics modules simulate RSVP-TE signaling in order to test DUT MPLS capabilities and features.
- Supports IP and UDP protocol in transit LSR and egress LER testing.
  Supports IP, TCP, UDP, and ICMP protocols in ingress LER testing.
- Tests a mixture of static labeled and dynamic labeled MPLS traffic.
  TeraMetrics cards support both static and dynamic labels. TeraMetrics-based cards support static labels.
- Allows full control over TSpec and session attribute parameters for traffic engineering.
- Provides explicit routing with loose or strict hops.
- Allows aggregation of LSPs.

Figure 17-6 illustrates how SmartFlow can be used to test MPLS. The results are for a test that uses a combination of MPLS and non-MPLS traffic and tests various levels of QoS. The results show a graph of Latency test results.

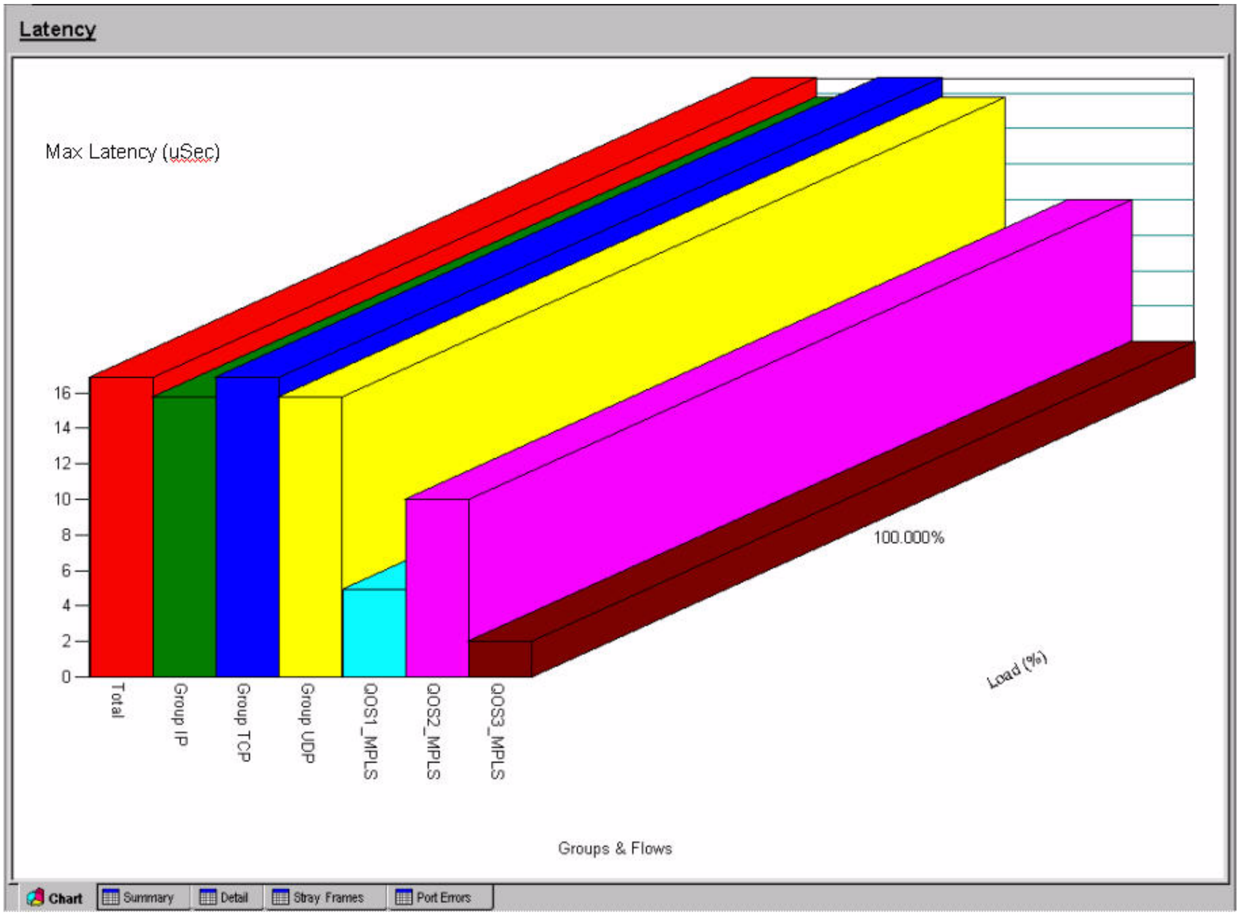

Figure 17-6. Latency Test Results

Installation Requirements for MPLS Testing in SmartFlow

This section lists the hardware and software requirements for QoS testing using the MPLS features of SmartFlow as well as TeraMetrics and TeraMetrics-based modules.

For information on SmartBits installation procedures, see the *SmartBits 600x/6000x Installation Guide*.

SmartBits Hardware

Use one of the following chassis:

- SmartBits 600x
- SmartBits 6000x.

You can use either chassis with any of the LAN or POS modules listed in *Table 3-3, "Supported TeraMetrics Modules," on page 62* and/or *"Supported TeraMetrics and TeraMetrics-based Modules" on page 62*. *Table 3-3* lists the TeraMetrics modules that can be used for both dynamic and static MPLS labels. *Table 3-4* lists the TeraMetrics-based modules that can be used for static MPLS labels.

To run the back-to-back transit LSR test, you must also have one RJ-45 cross-over cable with network connectivity to the chassis.

Note: TeraMetrics modules are required for dynamic MPLS labels. If you want to use static (non-negotiated) labels, use any TeraMetrics-based module; however, since there is no RSVP-TE signaling, the DUT needs to be configured statically.

Software

The following software is required to run MPLS in SmartFlow:

- The MPLS .rpm file (required for dynamic labels only)
- SmartFlow
- SmartBits Download Program Manager.
  (The Download Program Manager is packaged on its own CD.)

Installing the RPM File

MPLS functionality must be added to TeraMetrics modules by loading MPLS `.rpm` files using download manager. (The term "MPLS `.rpm`" is used generically. Refer to the Release Notes for the name of the current MPLS `.rpm` file).

Note: You must use the download manager to download firmware or RPM files to any TeraMetrics card that will be used for MPLS testing. Refer to the Release Notes for the card firmware or to the download manager online Help for details on using the download manager.

General Test Setup for MPLS Testing

Setting up a test for MPLS testing includes all of the steps needed to set up any other test. For general information about how to set up a test, see *Chapter 8, "Set up and Run Tests."*

These are additional steps relating specifically to MPLS testing that must be done:

1 If this is the first time that TeraMetrics cards are being used for MPLS testing in SmartFlow, install the MPLS .rpm file on those modules that will be used in MPLS testing.

See *"Installing the RPM File" on page 423* for details.

2 On the *IPv4 Networks* tab, select the **RSVP-TE** checkbox to enable MPLS signaling on the port.

The RSVP-TE signaling protocol is the means by which applications signal the QoS requirements to the network. The network responds with success or failure messages.

If testing an ingress LER, enable this field for the receiving port(s).
If testing an egress LER, enable this field for the transmitting port(s).
If testing a transit LSR, enable this field for both transmitting and receiving ports.

Notes: •

- RSVP-TE is only available for ports on TeraMetrics cards. The *RSVP-TE* field is grayed out for other cards, including TeraMetric-based cards.
- *For TeraMetrics cards*: If dynamic labels will be used, select the *RSVP-TE* checkbox for either the transmitting port, receiving port, or both ports, depending on the type of MPLS test. Note that the *RSVP-TE* checkbox can remain selected for both transmitting and receiving ports for an ingress LER, an egress LER, or for non-MPLS testing.

3 To test PHP with implicit null, select the **PHP** checkbox for the port on the *IPv4 Networks* tab.

SmartFlow supports PHP with implicit null, which always returns a value of **3**. [Select the **PHP** checkbox to test whether a transit LSR near the end of the tunnel "pops" (removes) the routed label in packets in an MPLS network and then forwards the unlabeled packets to the egress LER.]

4 Enter the IP address for the MPLS neighbor in the **MPLS Neighbor** field. This address is usually the same address as the IP port in the gateway (which is usually the router ingress port).

5 On the **SmartFlows>Traffic** tab, select the **Enable MPLS** checkbox.

To enable MPLS for multiple flows:
Highlight the flows, right-click, and select the **Enable MPLS** checkbox.

This checkbox is only enabled if the *RSVP-TE* checkbox on the *IPv4 Networks* tab is selected for the transmitting port.

6 Select either **IP** or **UDP** in the **IP's next protocol** field.

7 On the *MPLS LSP* tab, verify whether you want to use the default LSP or set up any additional LSPs.

8. On the **SmartFlows>MPLS** tab, assign (or verify an assignment of) either a static label or an LSP (specified at the **MPLS LSP** tab) to the flow.

   Specify whether the label should be dynamic or static in the **Dynamic Label** field.

   - If it is a static label, enter the label number in the **Label** field.
   - If it is a dynamic label, select the LSP that you want the flow to use. Only LSPs that were already set up at the *MPLS LSP* tab and whose *Ingress Port* value matches the transmitting port for the flow appear in the drop-down list.

   For more information, see *"Assigning Labels to Flows" on page 428*.

9. On the **SmartFlows>MPLS** tab, add more labels for the flow, if needed. Specify the parameters for the label as you did in the previous step.
10. On the **Options** tab, verify or change the amount of time that the DUT should wait for paths to be set up before timing out in the **LSP setup timer (Sec)** field in the **MPLS testing** pane. Select the **Allow retry** checkbox if you want the DUT to retry the setup if paths are not set up within the designated time.
11. Run the test.
12. For ingress LER and transit LSR tests, verify the labels negotiated for each flow in the `MPLSLabel.txt` file.

Enabling Ports for MPLS Signaling

The RSVP-TE signaling protocol is the means by which applications signal the QoS requirements to the network. The network responds with success or failure messages. Enable ports for MPLS signaling on the *IPv4 Networks* tab (as shown in *Figure 17-7*).

The following fields on the *IPv4 Networks* tab apply to MPLS testing:

- *RSVP-TE*
- *PHP*
- *MPLS Neighbor.*

For more information about the fields on the *IPv4 Networks* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar. The online Help also contains more details on enabling ports for MPLS.

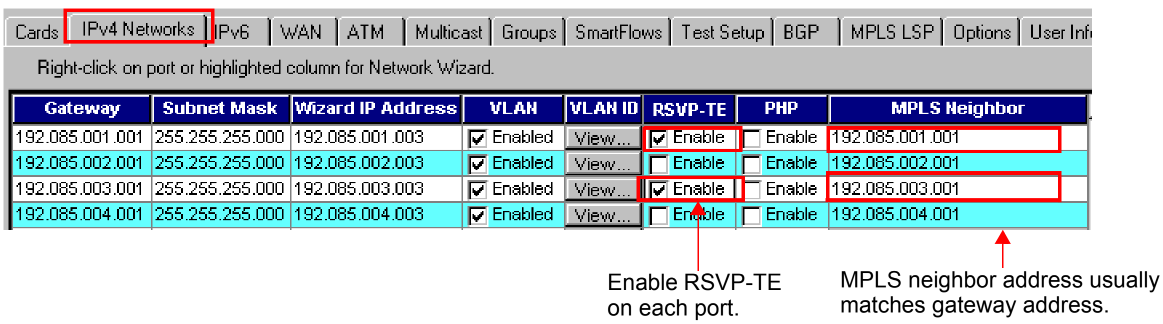

Figure 17-7. IPv4 Networks Tab

In addition to QoS requirements, the RSVP signaling protocol also carries traffic specifications such as controlled load that are set up on the *Setup TSpec* dialog box. For more information, see *"Setting Up a Controlled Load (TSpec)" on page 432*.

Setting Up and Adding LSPs

Once the MPLS-related fields have been configured on the *IPv4 Networks* tab (including selecting the appropriate *RSVP-TE* fields), set up more LSPs (in addition to the default LSP) on the *MPLS LSP* tab. *Figure 17-8 on page 427* illustrates the *MPLS LSP* tab.

Use the *MPLS LSP* tab to:

- Change attributes for the default LSP, such as the TSpec.
- Aggregate multiple flows for the same LSP, in which case an LSP must be added.
- Perform an ingress LER or egress LER test, which requires the addition of a dynamic LSP.

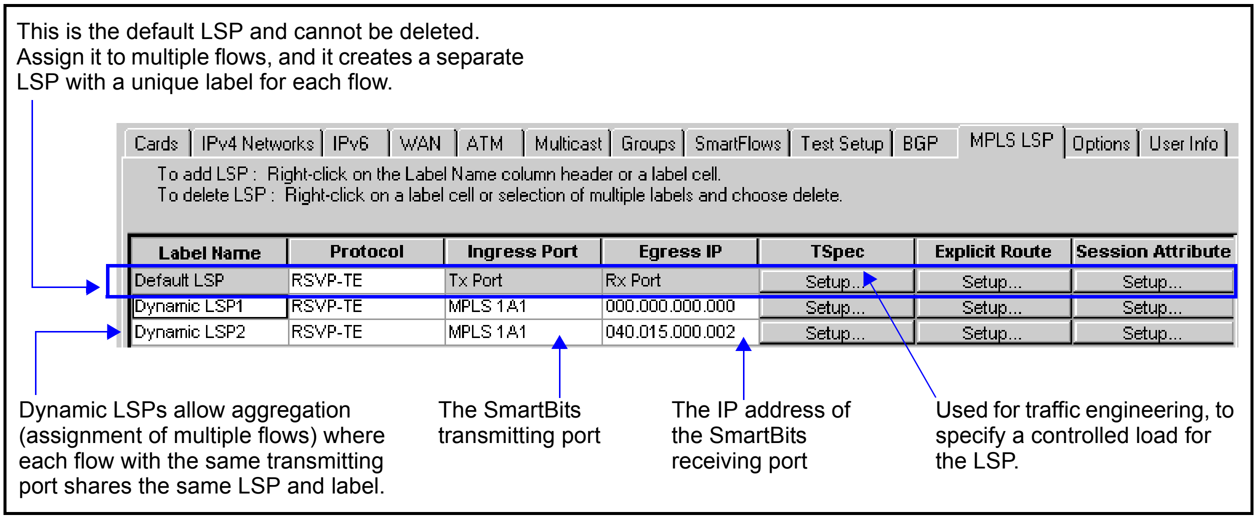

Figure 17-8. MPLS LSP Tab

The default LSP can be used without any modifcation.

Notice that the default LSP is listed first. Unlike the dynamic LSPs, you can assign more than one flow to the default LSP and each flow still gets a separate label. The default LSP value for the *Ingress Port* field is Tx Port, which is the SmartBits transmitting port. Notice also that the value for the *Egress IP* field is Rx Port**,** which is the address in the *Port IP Address* field of the SmartBits receiving port.

Notes:

- The *MPLS LSP* tab does not apply to ingress LER tests.
- For egress LER testing, use dynamic LSPs, not the default LSP.
- The DUT slows down as the number of LSPs increases.

All added LSPs are dynamic. Consequently, multiple flows can be assigned that will use the same LSP. Add a dynamic LSP if:

- You want multiple flows (with the same transmitting port) to use the same LSP and label.
- You want to set up multiple LSPs, each with the same TSpec attributes.

If LSPs are not added, all flows are automatically assigned to the default LSP. For information about how to assign flows to LSPs, see *"Assigning Labels to Flows" on page 428*.

For a description of each field on the *MPLS LSP* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

How the Default LSP Works

The default LSP is a special LSP because it can be used as a template to quickly and easily create multiple LSPs for multiple flows on one or more ports without having to modify anything for each LSP. Although it is a non-aggregated LSP, multiple flows can stll be assigned to it because it creates another LSP with a separate label for each assigned flow. However, the LSP name for each flow assigned to it is not visible in the user interface.

The default LSP automatically uses the transmitting port of the flow assigned to it as its *Ingress Port* value and the receiving port of the flow as its *Egress IP* value.

The LSP *Ingress Port* value should be the same as the IP address for the SmartBits transmitting port. The *Egress IP* value is the address in the *Port IP Address* field of SmartBits receiving port. Because MPLS is enabled on the transmitting port and the receiving port, a transit LSR test is executed.

Assigning Labels to Flows

Once LSPs are set up, assign these LSPs to flows transmitted from the same port on the *SmartFlows>MPLS* tab (*Figure 17-9 on page 428*). Use this tab to specify:

- Multiple labels for the flow.
  Each packet of the flow has multiple MPLS label stack entries (shims). Multiple labels per packet allows you to consolidate (encapsulate) multiple LSPs as a way to identify them as all from one source, such as a service provider.
- Static verus dynamic labels.
  Either the label number (static) is specified or the DUT negotiates it (dynamic). (TeraMetrics-based cards only support static labels.) Use static labels as a diagnostic tool on an MPLS network (if RSVP-TE is disabled) or if you simply know the path that you want the label to take.
- TTL label fields.
  First, set the experimental bit of the MPLS label stack.

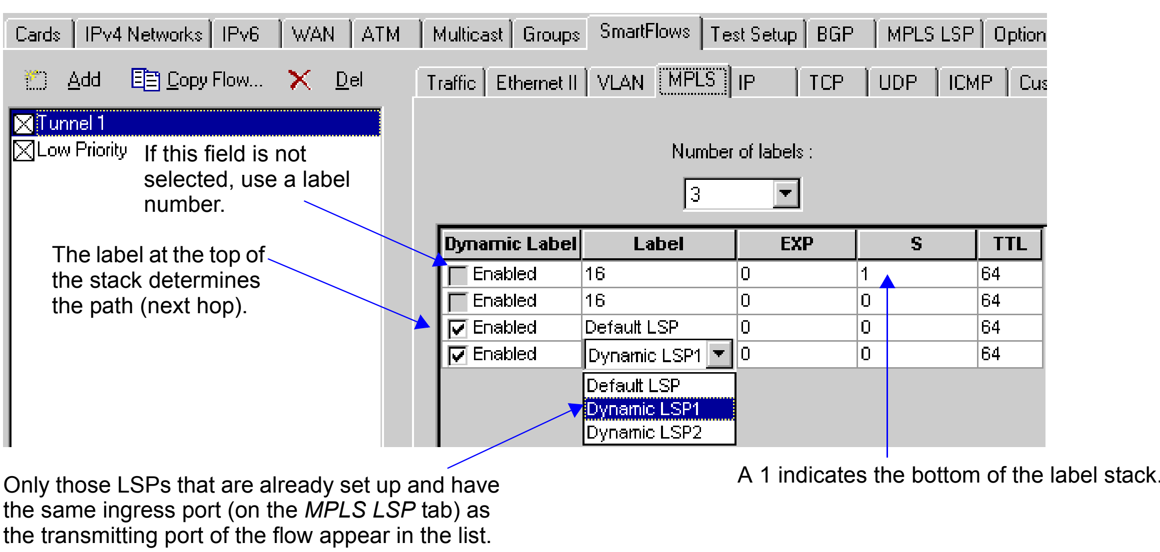

Figure 17-9. SmartFlows>MPLS Tab

Tip: If none of the fields on this tab can be edited, verify that the *Enable MPLS* checkbox on the *SmartFlows>Traffic* tab is selected.

For a description of each field on the *SmartFlows>MPLS* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar. The online Help also contains more detailed information about assigning labels to flows.

Notes:

- When static labels are used, use a label number instead of assigning an LSP.
- If an LSP does not appear in the *Label* list that was set up, make sure the ingress port on the *SmartFlows>MPLS* tab matches the transmitting port for the flow that you are trying to assign to an LSP. *Figure 17-10 on page 430* shows this relationship.

Aggregated LSPs (Multiple Flows per LSP)

An aggregated LSP is a dynamic LSP with multiple flows assigned to it. Aggregate LSPs when you want:

- To create a large amount of traffic for the LSP in order to make it more realistic.
- To simulate and test traffic engineering, where all flows of a certain type (such as video) use the same LSP.

Flows can be assigned to the same dynamic LSP only if they have the same transmitting port. (The *Ingress Port* field value matches the transmitting port of the flows.)

Figure 17-10 illustrates the relationship between the value in the *Ingress Port* field on the *MPLS LSP* tab, the transmitting port for the flow, and the LSPs that display for dynamic labels.

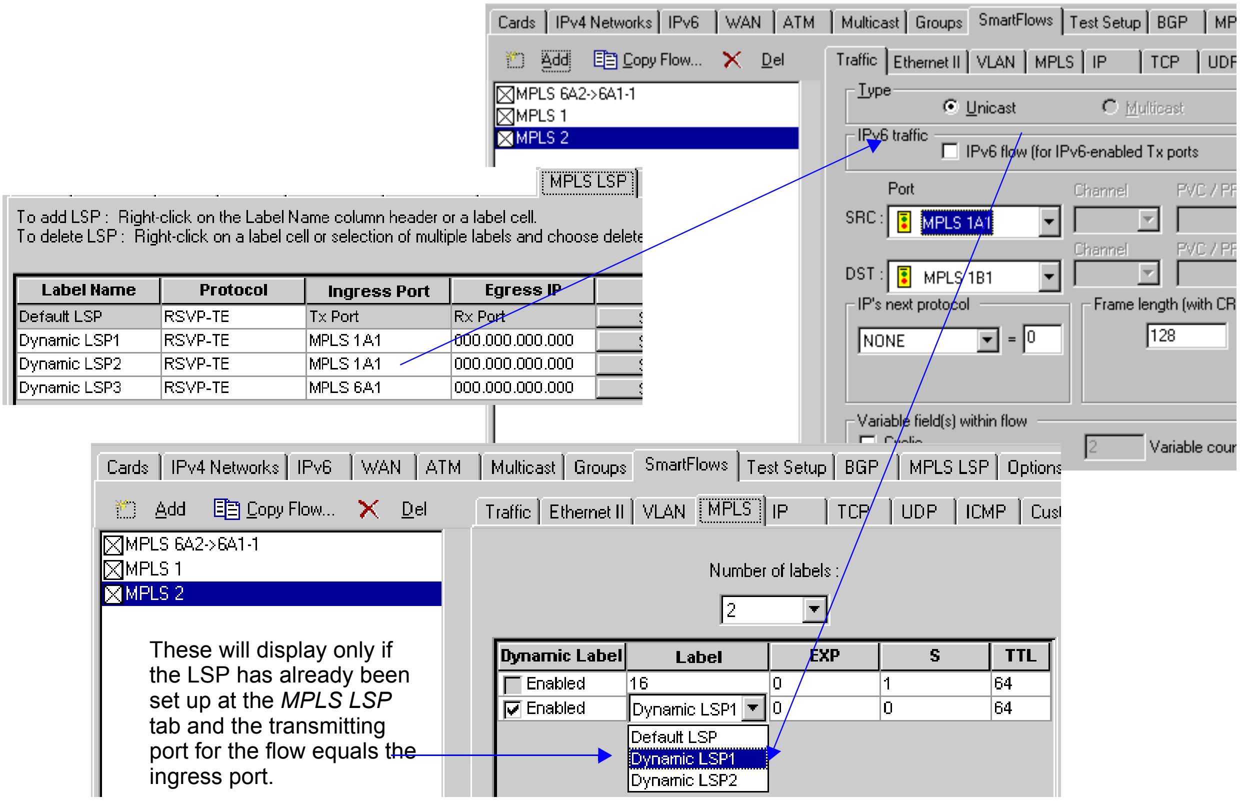

Figure 17-10. How to Aggregate an LSP at the SmartFlows›MPLS Tab

To aggregate LSPs, from the *SmartFlows>MPLS* tab, select the same dynamic label in the **Label** field for the flows that you want to aggregate.

When you aggregate LSPs, the `mplslabels.txt` file will only contain one label, which is the label used on the transmitting SmartBits port.

Verifying Flow Labels

Once the test is complete, check the `MPLSLabel.txt` file (*Figure 17-11*) to see if the labels were set up correctly.

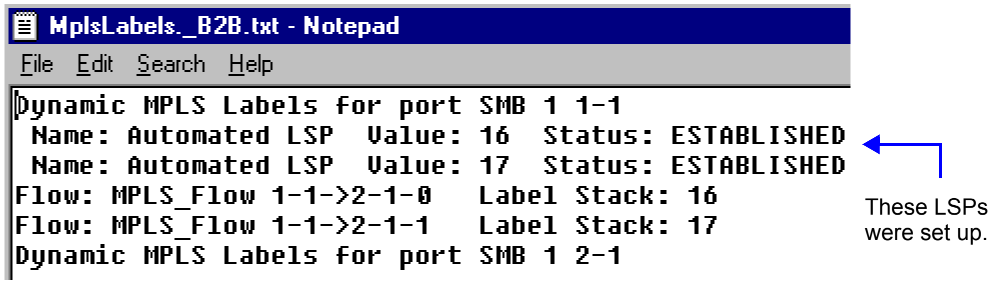

Figure 17-11. Sample MPLSLabel.txt File

For the procedure to do this, refer to the online Help.

Flow Labels for Aggregated LSPs

If multiple flows are assigned to one LSP (aggregated), the `MPLSLabel.txt` file should reflect this fact. The flows that were assigned to the same LSP should have the same label number. Suppose there are four flows. Flows 0 and 1 are assigned to different LSPs, and Flows 2 and 3 are assigned to the same LSP. The `MPLSLabel.txt` file could look like the file shown in *Figure 17-12*. In this file, Flows 0 and 1 negotiated different labels, and Flows 2 and 3 share the same label.

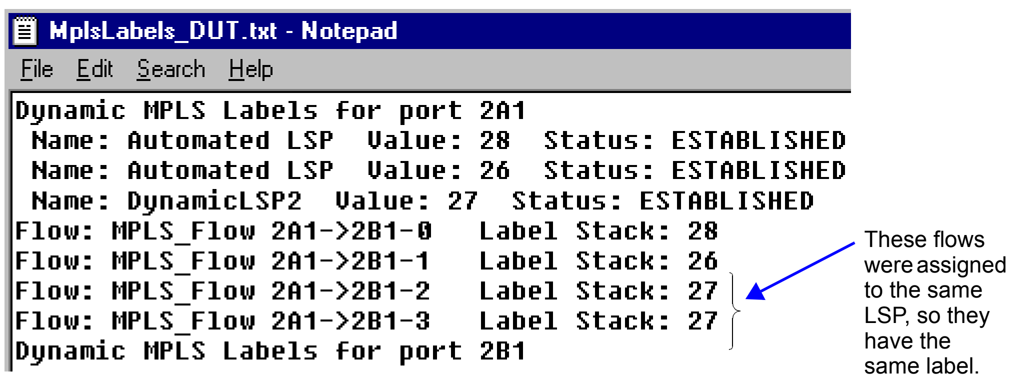

Figure 17-12. Sample MPLSLabel.txt file with Aggregated LSPs

Setting Up LSPs with Traffic Engineering

When an LSP is set up, use the default values for the TSpec and session attributes, or modify these attributes to more closely align with traffic engineering goals.

Both the Tspec and session attributes parameters help define the traffic that should flow through the LSP.

Note: For the signaling phase, there is no guarantee that the DUT is capable of the parameters that are specified in the *Setup TSpec* or *Session Attribute* dialog boxes.

Setting Up a Controlled Load (TSpec)

The *Setup TSpec* dialog box (*Figure 17-13*) contains parameters that describe the traffic that will be put onto the network or the traffic an application expects to generate at the time the tunnel is being set up. These parameters apply to the controlled load element of a QoS control service. They define attributes such as rate and frame size requirements that frames must meet in order to use a particular LSP.

The controlled load element employs the concept of a token bucket. A token bucket is a traffic-shaping mechanism in which a flow's ability to send frames is controlled by the presence of tokens. A token represents a unit of bytes.

To access the **Setup TSpec** dialog box from the **MPLS LSP** tab, click the **Setup** button in the **TSpec** column for the appropriate LSP.

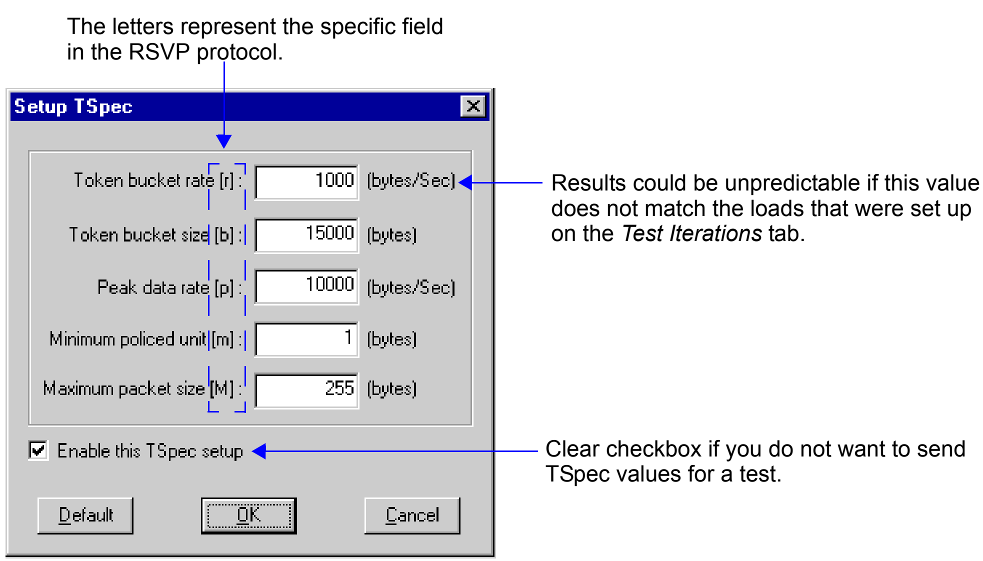

Figure 17-13. Setup TSpec Dialog Box with Default Values

Notes:
- If you do not select the **Enable this TSpec setup** checkbox, no Tspec values are sent. Some devices require TSpec parameters. Consult the appropriate vendor documentation to determine whether or not the setup is needed.
- Anything sent in these frames that violates these parameters may cause those frames to be discarded. These parameters should be consistent with the frame size(s) and rates specified in the test setup.

For a description of each field on the *Setup Tspec* dialog box, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Setting Up Session Attributes

The *Session Attribute* dialog box (*Figure 17-14*) contains parameters that can be added to RSVP path messages to help identify the session for diagnostic purposes. These parameters apply to the MPLS attribute object. They define attributes such as priority and rerouting.

To access the **Session Attribute** dialog box from the **MPLS LSP** tab, click the **Setup** button in the **Session Attribute** column for the appropriate LSP.

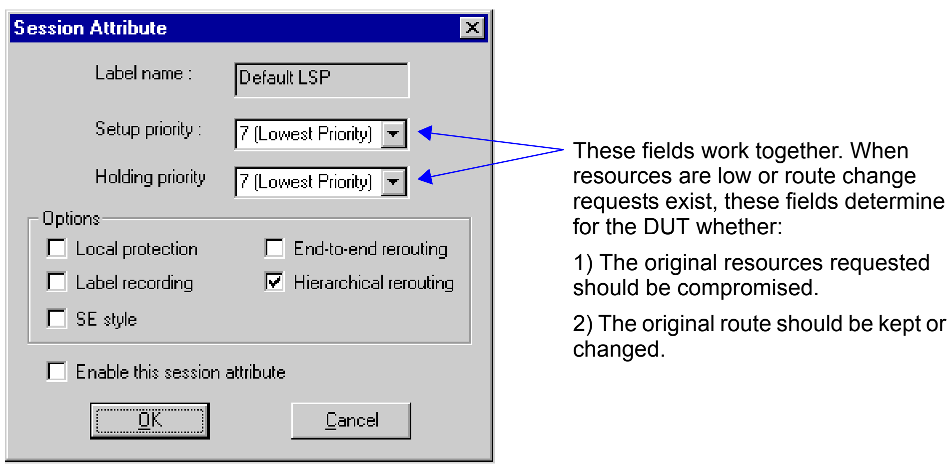

Figure 17-14. Session Attribute Dialog Box

For a description of each field in the *Session Attribute* dialog box, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

How Setup Priority and Holding Priority Work Together

When either a route change is requested or resources are low, both the *Setup priority* and *Holding priority* fields determine for the DUT what priorities the LSP should receive.

The recommended value for both fields is 4, so that neither route selection nor resource selection is dominant.

Example 1 - Higher Setup Priority

Setup priority: 5
Holding priority: 3

Since the setup priority is higher than the holding priority, the DUT keeps the resource requirements originally requested for the LSR rather than negotiate for resources even if they are low. However, since the holding priority is lower, it may negotiate a different path.

Example 2 - Higher Holding Priority

Setup priority: 3
Holding priority: 5

Since the holding priority is higher than the setup priority, the DUT keeps the original route even if a route change is requested or resources are low rather than negotiate for a different path. However, since the setup priority is lower, it may negotiate or compromise the original resources requested for the LSR.

Setting Up Explicit Routes

An explicit route is one that is specified as a sequence of hops rather than being determined solely by conventional routing algorithms on a hop-by-hop basis. It is a static, rather than negotiated route. The explicit route object is used to specify the route that RSVP path messages take for setting up LSPs. TeraMetrics-based cards only support explicit routes.

Explicitly routed LSPs allow traffic to be engineered to travel through the network, avoiding congestion, and according to QoS requirement, security classification, etc. The capability to do explicit routing (LSPs) within MPLS helps the network manager engineer a network that sustains QoS. LSPs can be manually created through the network to ensure QoS guarantees.

You can edit the explicit routes attribute for both the default LSP and dynamic LSPs on the *Setup Hop* dialog box.

An explicit route can be defined either of two ways:

- *Strict* – Specifies and enforces the next directly connected router in the LSP. Select the **Strict** checkbox on the **Setup Hop** dialog box to use this option.

  This option is useful if using the controlled load (TSpec) QoS service control to ensure that the LSP uses routers with enough bandwidth to support the TSpec parameters. You can also use this option if, for security reasons, you require an LSP to be established only over specific, trusted routers.

- *Loose* – Relies on the routing table to find a destination. Clear the **Strict** checkbox on the **Setup Hop** dialog box to use this option.

  The option allows the router to use an alternate route, if needed. The router may change the route for various reasons, including:

  - The specified explicit route is torn down.

- A specified router along the LSP is down.
- The specified explicit route is congested, so a route with less traffic is more desirable.
- The specified router cannot handle any more traffic.

You can edit the explicit routes attribute for both the default LSP and dynamic LSPs on the *Setup Hop* dialog box. *Figure 17-15* shows five (the maximum allowed) explicit routes, with four strict routes (where the *Enabled* checkbox is selected) and one loose route (where the *Enabled* checkbox is clear).

To access the **Setup Hop** dialog box from the **MPLS LSP** tab, click the **Setup** button in the **Explicit Route** column for the appropriate LSP.

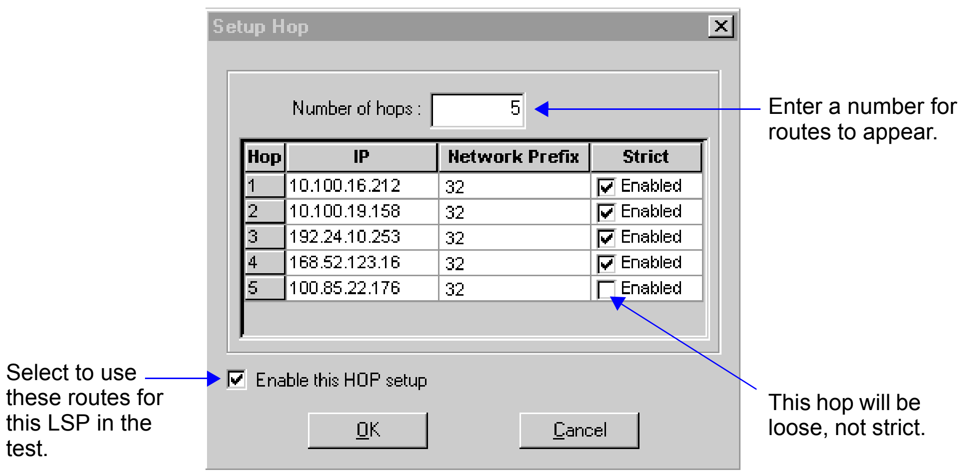

Figure 17-15. Setting Up Explicit Routes

For a procedure on how to set up explicit routes for an LSP, refer to the online Help.

Testing MPLS

If MPLS is the key to large network management, LSR testing is the key to successful MPLS deployment. To properly assess MPLS performance, it is necessary to test the following:

- The ability of the LSRs to set up LSPs
- The ability of the LSRs to swap labels
- The QoS functionality of the DUT when deployed as:
    - An ingress LER (Does it accurately classify packets, apply labels, and forward packets?)
    - A transit LSR (Does it accurately swap and forward packets?)
    - An egress LER (Does it remove the label and forward the packet correctly?)

Test Configurations

To test the capabilities and features of an MPLS-enabled device in each of its possible roles, three test scenarios are recommended:

- Ingress LER test
- Transit LSR test
- Egress LSR test.

The type of MPLS test scenario is determined by:

- Which ports (transmitting, receiving, or both) are enabled in the *RSVP-TE* field (on the *IPv4 Networks* tab).
- The values in the *Ingress Port* and *Egress IP* fields (on the *MPLS LSP* tab).
- Which flows are MPLS-enabled (on the *SmartFlows>Traffic* tab).

Diagrams showing the respective roles of the SmartBits modules and of the DUT in each of these test scenarios will help you correctly set up the type of MPLS test that is needed. See *Figure 17-16 on page 437*, *Figure 17-17 on page 437*, *Figure 17-18 on page 438*, *Figure 17-19 on page 439*, and *Figure 17-20 on page 439*.

Ingress LER Test

Figure 17-16 illustrates the test setup for a transit LSR test.

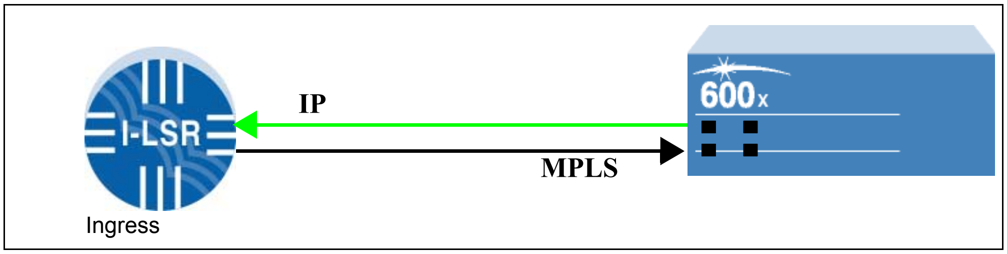

Figure 17-16. Ingress LER and SmartBits Modules

Purpose: The Ingress LER test allows you to ensure that a DUT can receive IP or other non-MPLS traffic and successfully insert an MPLS label into the packets.

During an Ingress LER test, the following occurs:

1. An LSP is set up between the DUT and the receiving SmartBits port.
2. The SmartBits transmitting port sends IP packets to the DUT (ingress LER).
3. The DUT inserts MPLS labels into the packets and sends them to the SmartBits receiving port.

The *MPLS LSP* tab is not applicable when an ingress LER test is set up. When an ingress LER test is run, the receiving ports are configured with RSVP-TE (on the *IPv4 Networks* tab). Signaling occurs between the DUT and the receiving ports.

Important: Before running the ingress LER test, enter a value in the *Delay before transmit (Sec)* field in the *Global test options* pane on the *Options* tab to allow sufficient time for MPLS signaling between the DUT and the SMB receiving port to complete. (The recommended value is 50 seconds.) This time delay prevents test traffic from being sent before signaling is complete.

For a detailed procedure on how to run the Ingress LER test, refer to the online Help.

Transit LSR Test

Figure 17-17 illustrates the test setup for a transit LSR test.

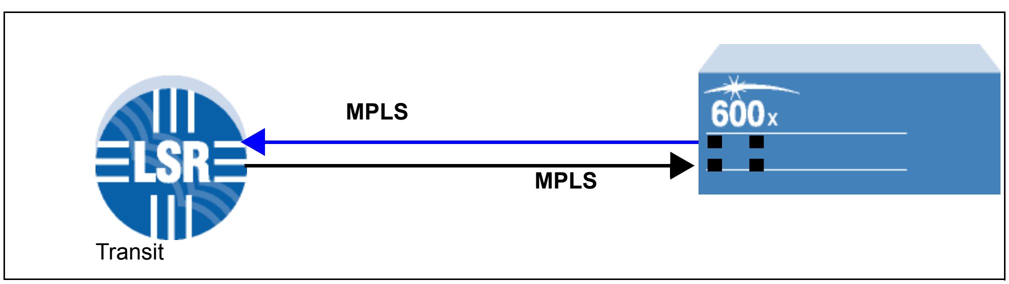

Figure 17-17. Transit LSR and SmartBits

During a transit LSR test, the following occurs:

1. An LSP is set up between the SmartBits transmitting port and the SmartBits receiving port.
2. The SmartBits transmitting port (acting as ingress LER) sends MPLS packets to the DUT (transit LSR).
3. The DUT swaps labels and sends MPLS packets to the SmartBits receiving port (acting as the egress LER).

Two test configuration methods for a transit LSR test are as follows:

- Back-to-back configuration (test setup as a practice exercise)
- SmartBits with DUT.

Back-to-Back Transit LSR Test

Purpose: This test is more of an exercise designed to introduce you to test setup basics and to demonstrate MPLS functionality on the SmartBits equipment. If necessary, it can also be used to verify equipment functionality.

Figure 17-18 shows a back-to-back transit LSR test configuration using a SmartBits 600x with two LAN-3301A modules.

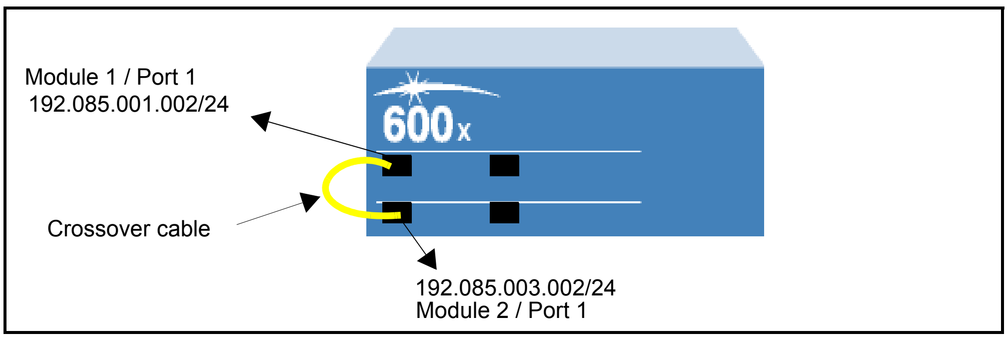

Figure 17-18. Back-to-back Transit Test Configuration

For a detailed procedure on how to run the back-to back transit LSR test, refer to the online Help.

Transit LSR Test with DUT

Purpose: The purpose of this test is to determine SmartFlow test results when transmitting MPLS traffic through a DUT.

Figure 17-19 shows a transit LSR test configuration using a SmartBits 600x, two LAN-3301A modules, and a DUT. Use these sample addresses to run the MPLS test or enter your own addresses.

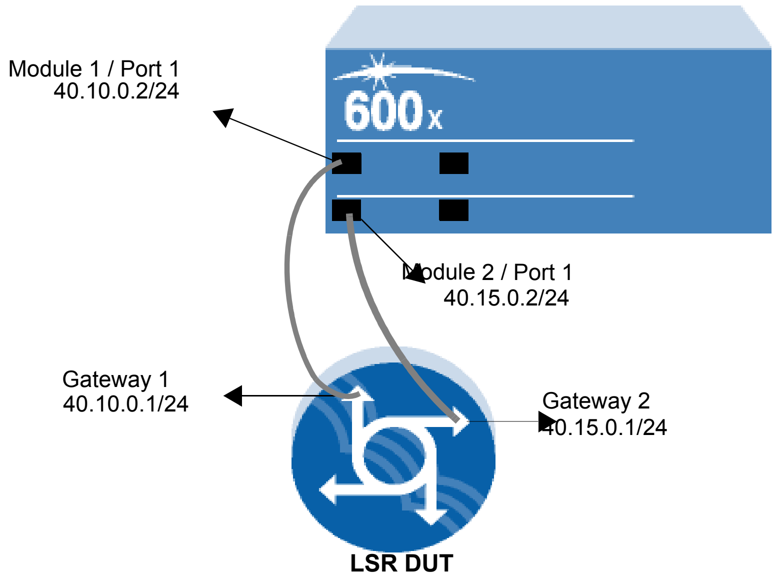

Figure 17-19. Sample Transit LSR Test Configuration (with DUT)

For a detailed procedure on how to run the transit LSR with DUT test, refer to the online Help.

Egress LER Test

Figure 17-20 illustrates the test setup for an egress LER test.

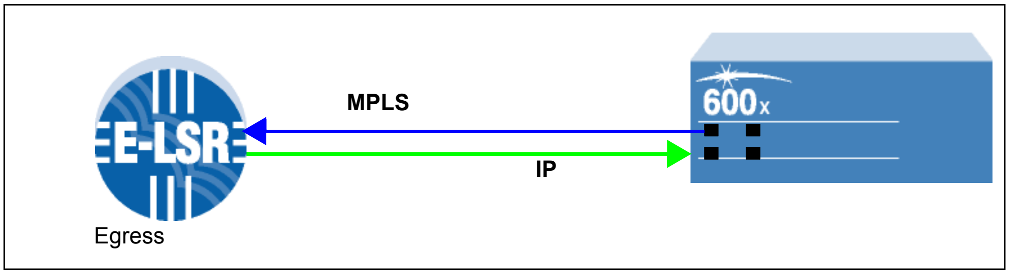

Figure 17-20. Egress LER and SmartBits

Purpose: The egress LER test allows you to ensure that a DUT can receive MPLS traffic and successfully remove the MPLS label from the packets and forward IP or other non-MPLS traffic.

During an egress LER test, the following occurs:

1 An LSP is set up between the SmartBits transmitting port and the DUT.

2 The SmartBits transmitting port (acting as ingress LER) sends MPLS packets to DUT (egress LER).

3 The DUT pops labels and sends IP packets to the SmartBits receiving port.

Important: Before running the egress LER test, enter a value in the *LSP setup timer (Sec)* field in the *MPLS testing* pane on the *Options* tab to allow sufficient time for MPLS signaling between the DUT and the SmartBits transmitting port to complete. (The recommended value is 50 seconds.) This time delay prevents test traffic from being sent before signaling is complete.

For a detailed procedure on how to run the egress LER test, refer to the online Help.

Chapter 18

Testing BGP Routers

This chapter contains information about the BGP4 advertisement features of SmartFlow and general information about BGP technology. These capabilities allow you to generate and advertise a large number of BGP4 routes.

In this chapter...

Overview of BGP

Border Gateway Protocol (BGP) is a routing protocol used in Internet networks that are based on Internet Protocol (IP). A routing protocol is defined by a set of message formats that describe the reachability and preference of a variety of network addresses. The protocol also includes rules for processing the information that is received through its messages. The basic role that a routing protocol plays is to ensure that information can be sent between computers that are connected to the network. BGP is used exclusively to connect various networks; hence it is classified as an *inter-domain* routing protocol.

Routing is the complex part of data communication between networks. Once a route is determined, the transmission of information groups (packets) is rather straightforward. Determining the path to take to transport those packets is another issue. The BGP protocol addresses the task of path determination in networks. It is used to exchange routing information between gateway hosts in a network of autonomous systems.

BGP performs inter-domain routing in Transmission-Control Protocol/Internet Protocol (TCP/IP) networks. BGP is an Exterior Gateway Protocol (EGP), meaning that it performs routing between multiple autonomous systems or domains, and it exchanges routing and reachability information with other BGP systems. (See *Figure 18-1 on page 443*.)

The following are some of the primary characteristics of BGP:

- BGP supports multiple connections between autonomous systems and provides for both primary and backup links.
- Routing information includes destination networks, next hop destinations, and autonomous system number lists.
- BGP Version 4 (BGP4) includes support for Classless Inter-Domain Routing (CIDR).
- Each core BGP router knows routes to every defined network in the router networking environment.

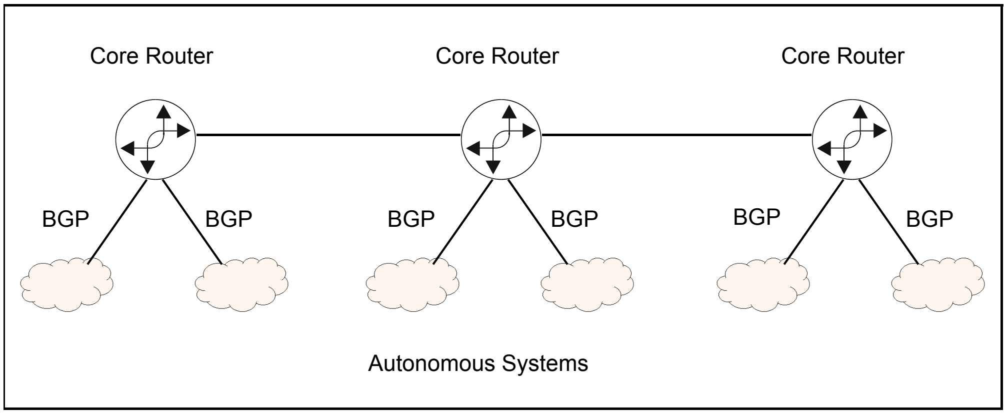

Figure 18-1. Core Routers Using BGP for Traffic Routing Between Autonomous Systems[1]

Autonomous Systems and Routing Protocol Types

An Autonomous System (AS) is a collection of networks under the control of one administrative authority (e.g., UUNET, Qwest, MichNet, etc.). The assignment of autonomous systems is strictly controlled.

There are various types of routing protocols, of which BGP is one type. Interior protocols run only within one AS. These are typically "enterprise" or ISP protocols such as OSPF, RIP, and IS-IS. Exterior protocols are typically "core" network protocols and include BGP and EGP. These protocols run between different autonomous systems.

A Replacement for EGP

BGP was developed to replace its predecessor, the now obsolete Exterior Gateway Protocol (EGP), as the standard exterior gateway-routing protocol used in the global Internet. The advent of BGP solved some serious problems associated with EGP. BGP also scales better than EGP, adapting itself more efficiently to Internet growth.

BGP is specified in the following Request for Comments (RFCs):

- RFC 1771: Describes BGP4, the current version of BGP.
- RFC 1654: Describes the first BGP4 specification.
- RFC 1105, RFC 1163, and RFC 1267: Describe versions of BGP prior to BGP4.

BGP Routing

BGP is the protocol that is usually used between gateway hosts on the Internet. As with any routing protocol, BGP maintains routing tables, transmits routing updates, and bases

---

1. Source: *Border Gateway Protocol (BGP),* Cisco Press, 1999.

routing decisions on routing metrics. The primary function of a BGP system is to exchange network-reachability information, including information about the list of autonomous system paths, with other BGP systems. This information can be used to construct a graph of autonomous system connectivity from which routing loops can be pruned and with which autonomous system-level policy decisions can be enforced.

Each BGP router maintains a routing table that lists all feasible paths to a network. This includes a list of known routers, the addresses they can reach, and a cost associated with each path to each router. All this information enables the best route to be chosen for a given data session. The router does not refresh the routing table. Instead, routing information received from peer routers is retained until an incremental update is received.

BGP devices exchange routing information upon initial data exchange and after incremental updates. When a router first connects to the network, BGP routers exchange their entire BGP routing tables. Similarly, when the routing table changes, routers send the portion of their routing table that has changed. BGP routers do not send regularly scheduled routing updates, and BGP routing updates advertise only the optimal path to a network.[1]

Why Route Testing is Important

A BGP router (*Figure 18-2*) contains a switching fabric and multiple ports. The router is connected to a central CPU, a central route table, and additional autonomous systems.

The central routing table is a shared resource for all ports contained in the autonomous systems. Each BGP router port contains both route and ARP caches.

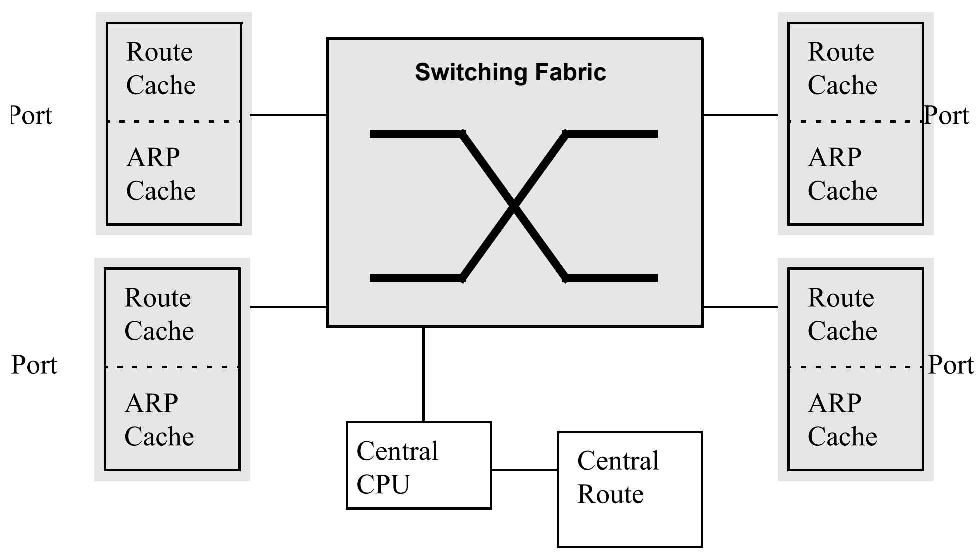

Figure 18-2. BGP Router

1. Source: *Border Gateway Protocol (BGP),* Cisco Press, 1999.

The BGP protocol is typically used for communication between core Internet routers. Unlike edge routers, which typically have a large ARP cache and a small route cache, core routers contain a small ARP cache and a large route cache. A typical BGP core Internet router contains more than 70,000 routes. To determine the proper output port, these routers must perform prefix matching on each and every packet. Testing this prefix matching ensures that the best possible routes are being selected.

Routing performance is determined by the performance of both the route cache and the switching fabric. This means that overall performance depends on the shared central route table resource and on individual port routing. The BGP feature of SmartFlow allows you to test the routing capabilities of the individual ports.

Proper Packet Routing

Packet routing generally follows the following steps:

1. Packets are stripped of their Layer 2 protocol.
2. Packets are classified at the ingress interface.
3. Packets are tagged with an output port based on route information. The ingress interface route cache is checked first. Next, the central route table is checked and the cache is refreshed.
4. Packets are submitted to the switching fabric.
5. Packets are augmented with the Layer 2 protocol of the egress port. Generally, all Layer 2 information is known because a core router is usually directly connected to another router. (Address Resolution Protocol [ARP] is usually not a factor.)

Identifying Poor Routing Performance

Routing performance is poor when route caches do not have necessary routes, when route caches are empty, or when route tables are too large.

BGP Route Flapping

When a router goes down, it causes a change in the routes (IP addresses of networks or groups of networks) stored in the BGP table on other routers. This causes the transmission of routing updates to advertise a destination network first via one route, then via a different route. This resulting change in routes is known as route flapping.

How a core router performs and handles priority traffic in the presence of route flapping is what differentiates it from other less superior routers. Ideal router performance with flapped routes consists of:

- Flapped routes that are forwarded (failed over) to a backup link.
- Unaffected routes that suffer no performance degradation.

Causes of Route Flapping

Possible causes of route flapping are:

- A bug in routing software
- A router going down (causing the link to go down)
- A problem with the circuit (causing the link to go down)
- Other network instabilities.

Why Route Flapping is a Problem

The amount of time route flapping takes depends on the number of routes. Route flapping typically takes 30 seconds between the disappearance (flap) and reappearance (unflap) of routes, but can also last for days until the basic problem with the router or circuit is fixed. A large number of route flaps at any one time is a serious problem. When this occurs in a BGP network, the amount of the routing information being advertised and withdrawn is considerable, and therefore greatly affects performance. BGP routers have to be configured to dampen (suppress), failover, or otherwise handle route flapping.

How Route Flapping Can Be Handled

There are several ways of handling route flapping. One method is called *dampening*. This method suppresses the advertisement of the route somewhat close to the route flapping source until the route becomes stable again. It stores a penalty value for each route and takes into account the past stability of a route (based on that penalty value) in deciding whether to use or readvertise that route.

Another method is *aggregation*, which can mask flapping. The disadvantage of using this method is that many sites use legacy addresses that cannot be aggregated.

What is the BGP Feature of SmartFlow?

The SmartFlow BGP feature provides the ability to generate and advertise a large number of BGP4 routes. SmartFlow supports only EBGP (external BGP) rules. The advertisement of routes is supported by the network interface card located in the PC on which SmartFlow is installed. Since SmartFlow advertises BGP routes before the MAC address learning phase, it puts the BGP route information in the Smrtflow.inifile so that the routes are in the router when SmartFlow sends ARP requests.

BGP Feature Capabilities

The SmartFlow BGP feature provides:

- The ability to advertise thousands of networks and subnetworks on all flow ports to a BGP4 router.

- The ability to advertise either on command or at regular preset intervals, and remove them on command or when the test is not running.
- The ability to set up the BGP session between the PC host and the router.
- The ability to flap and unflap network and/or additional routes once, continuously, or at specified intervals.

Caution: Due to limited memory on SmartBits cards and modules, you may not be able to create flows for every route. The Ethernet 10/100 Mbps cards currently support 1,000 flows per port. The SmartMetrics Gigabit cards (POS-6500 and POS-6502) support 8,000 flows per port. The SmartMetrics Gigabit card LAN-3201 supports 7,000 flows per port. It is important to be aware of the flows and routes that are being set up, so that the memory support of the installed cards is not exceeded. If you attempt to exceed the memory support of a card, a warning message appears.

Why Advertise Routes?

The BGP feature of SmartFlow allows the creation of multiple networks and subnetworks that serve as routes for transmitting packets. After creating these routes, the software "advertises" them to the utilized router. This gives the router the preferred selection of routes to use for transmitting packets. Flows can also be created which can pass between any of the created routes.

The advertisement of BGP routes allows the creation of flows between two external networks that are not directly connected to the router. The following example helps illustrate this situation.

To advertise BGP routes:

1 On the *IPV4 Networks* tab, set up the networks and subnetworks for Ports 1 and 2 of the SmartBits unit.

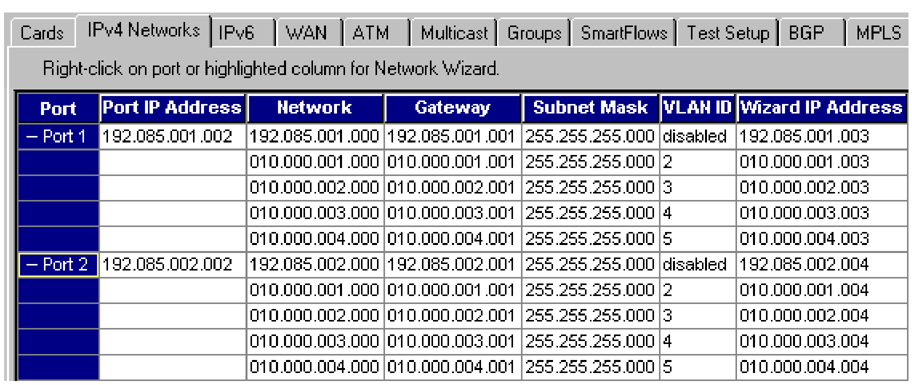

Port	Port IP Address	Network	Gateway	Subnet Mask	VLAN ID	Wizard IP Address
– Port 1	192.085.001.002	192.085.001.000	192.085.001.001	255.255.255.000	disabled	192.085.001.003
		010.000.001.000	010.000.001.001	255.255.255.000	2	010.000.001.003
		010.000.002.000	010.000.002.001	255.255.255.000	3	010.000.002.003
		010.000.003.000	010.000.003.001	255.255.255.000	4	010.000.003.003
		010.000.004.000	010.000.004.001	255.255.255.000	5	010.000.004.003
– Port 2	192.085.002.002	192.085.002.000	192.085.002.001	255.255.255.000	disabled	192.085.002.004
		010.000.001.000	010.000.001.001	255.255.255.000	2	010.000.001.004
		010.000.002.000	010.000.002.001	255.255.255.000	3	010.000.002.004
		010.000.003.000	010.000.003.001	255.255.255.000	4	010.000.003.004
		010.000.004.000	010.000.004.001	255.255.255.000	5	010.000.004.004

2 Set up a backbone mesh test to test the flows between the subnetworks.

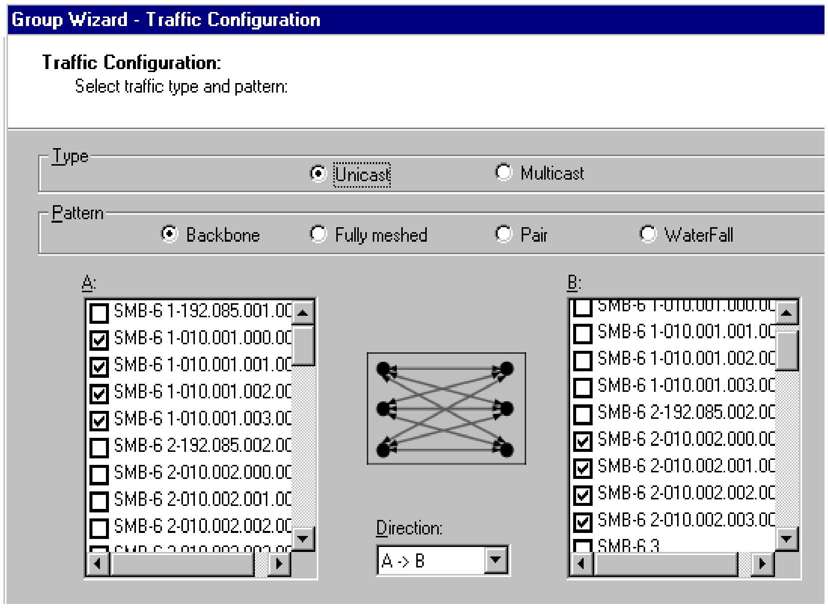

3 Assign values on the *BGP* tab, and then advertise the routes that were just set up. You also set parameters to flap and unflap routes on this tab.

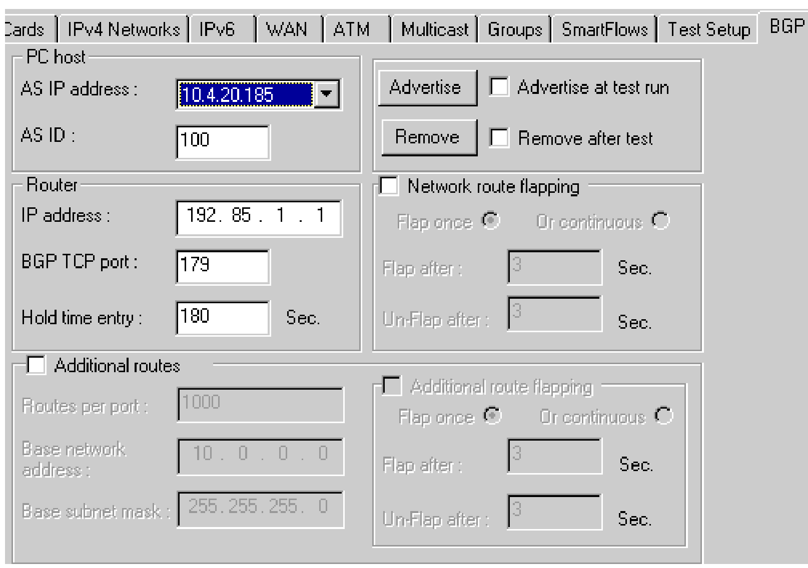

After running a test, the test results should indicate that some frames got through. This indicates that the advertisement of routes successfully informed the router of an available route to use. In this case, if routes were not advertised, the router would not know which route to use and test results would indicate 100% frame loss (i.e., the test would fail).

Basic BGP Testing Theory

SmartFlow BGP Operation

The BGP feature of SmartFlow includes:

- The generation of a BGP session.
- The advertisement of created routes to the BGP router.
- The removal and re-advertisement of created routes to the BGP router.

Route and Flow Creation

Routes are set up within the *IPv4 Networks* tab and additionally on the *BGP* tab. Each network can contain many subnetworks. Because very large numbers of groups and flows are needed to accurately test the capabilities of a router, the Group Wizard should be used to create groups and flows. The total number of flows created is limited by the type of card used in the SmartBits unit. If you attempt to create more flows than is allowed by the card, SmartFlow generates an error message.

PC-to-Router Communication

PC-to-router communication is established through the generation of a BGP neighbor session. Such a session is supported by a network interface card (NIC) installed in the PC. You have two basic setup options:

- The PC contains two NICs. One NIC supports the proxy script used to communicate with the router; the other NIC supports communication with the SmartBits unit. This is the recommended configuration.
- The PC (with one NIC) is connected to a single network that is connected both to the router and the SmartBits unit.

After the supportive hardware is in place, the router can then be prepared for communication by configuring the BGP endpoint (peer) on the router.

Note: In order to connect a BGP router to a PC, the router must have a 10/100 Ethernet port. If the router does not have this type of port, bridge or switch to the PC’s 10/100 Ethernet port.

BGP Peer Emulation

The intent of the SmartFlow BGP feature test equipment setup is to emulate a connection between two real BGP endpoints (BGP peers). Such a connection (called a BGP session) depends on a reliable protocol that provides an orderly delivery of messages. This protocol (TCP) is the same protocol used in the PC-to-router-to-SmartBits setup used for BGP testing.

A Real BGP Scenario

Once a connection is established between two BGP peers, TCP messages can be delivered reliably. In a real BGP session, the two BGP peers exchange routes with each other, sending the prefix of every route that each BGP peer wants the other to know. For example, if Router A has 50,000 prefixes and it is configured to send all of them to Router B, a large number of messages are sent, advertising all of these prefixes. After these initial messages have been exchanged, router A only informs router B of any changes made. The BGP protocol does not require that router A refresh information about any of these prefixes to router B.

A BGP session stays in an established state unless an error occurs, which closes the TCP connection, transitioning the BGP peers to an idle state. If the connection between BGP neighbors is broken (e.g., if one of the routers goes down or if there is no longer a viable path between them), the BGP connection times out and each end stops using the routing information received from the other.

A TCP connection can also be broken if no keepalive signals are received by a BGP peer from another BGP peer for a preset period of time. The sending of keepalive signals is not a function of TCP, but rather of BGP. A keepalive signal is sent from endpoint A to endpoint B, indicating that endpoint A is still up and running.

A BGP Test Scenario

In a BGP test scenario, the same basic rules apply as in a real BGP routing environment, except that routes are advertised by the PC loaded using the SmartFlow BGP feature (acting as a BGP peer) to the connected router (the other BGP peer). Routes are set up on the *IPv4 Networks* tab and/or the *BGP* tab, then advertised (via options in the *BGP* tab screen) by sending route prefixes to the router. The BGP session stays in an established state until the connection is broken. This can occur in a BGP test scenario if any of the following occurs:

- The BGP router goes down.
- SmartFlow removes BGP routes via the *BGP* tab screen in the application. This can be done manually or set up to occur automatically after a test is run.
- No keepalive signals are received by the router from the PC for the duration of the **Hold time entry** field value set on the *BGP* tab screen. (Default is set at 180 seconds.)

Routes are only valid while a connection is up.

BGP Route Advertisement

After setting up routes, groups, and flows, a user can advertise the routes to the connected router. To advertise the routes, the application first opens a connection to a BGP neighbor port. Next, for each SmartBits port, the application downloads the IP address of the port as the next hop. The next hop is the node to send data packets to in order to get the packets closer to the destination. In addition to the port IP address (next hop), the subnetworks assigned to that address are also downloaded.

This entire advertisement process provides the connected router with the available networks from which to choose routing packets. By loading each of these networks with thousands of flows, the routing capabilities of the router can be tested.

Advertisement Options

BGP routes can be advertised either manually, at regular intervals, or at the time a test is run. The latter option is recommended since the advertised routes are the most apt to represent reality (since changes could be made to routes or flows in between test runs).

When routes are advertised, they appear in the following format:

[Next Hop] [AS #] [Network 1, Network 2, Network 3...], where

Next Hop = the IP address of the SmartBits port.

AS # = the ID of the autonomous system to which the source network belongs (as defined on the *IPv4 Networks* and *BGP* tabs).

Network x = the address of the subnetworks assigned to the next hop.

An example of displayed route advertisement information appears below.

192.85.1.2	123	192.85.1.1/24,	10.1.1.0/24,	10.2.1.0/24,	10.3.1.0/24
192.85.2.2	123	192.85.2.1/24,	10.1.2.0/24,	10.2.2.0/24,	10.3.2.0/24
192.85.3.2	123	192.85.3.1/24,	10.1.3.0/24,	10.2.3.0/24,	10.3.3.0/24
			10.1.4.0/24,	10.2.4.0/24,	10.3.4.0/24

Recommended BGP Testing Strategy

To obtain the best representation of actual routing performance, it is recommended that a router be tested using a variety of route capacities. This approach provides a good idea of how effectively the router transmits different levels of flows from one port/network to another. For example:

- First run a test that employs only one route per port, with many flows on just a few of the routes.
- Next, run a test with one flow per route, using the maximum number of flows that a module can support. The ML-7710 cards currently support 1,000 flows per port. The SmartMetrics Gigabit cards (POS-6500 and POS-6502) support 8,000 flows per port. The SmartMetrics Gigabit card LAN-3201 supports 7,000 flows per port.
- Finally, run a test using about ten times the number of routes than flows, while using the maximum number of flows that a module can support.

PC Hardware Configurations

When testing BGP routes, two hardware configurations are possible:

- PC with one NIC
- PC with two NICs.

PC with One NIC

Figure 18-3 shows the required cabling setup for a PC containing one NIC.

You can also use the COM port to connect to a SmartBits 200/2000 chassis.

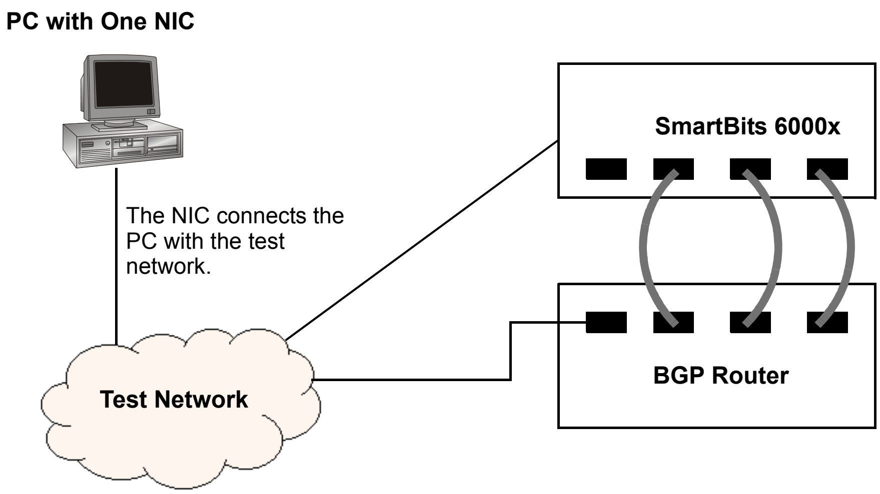

Figure 18-3. PC/Hardware Configuration for a PC with One NIC

PC with Two NICs

Figure 18-4 shows the required cabling setup for a PC containing two NICs.

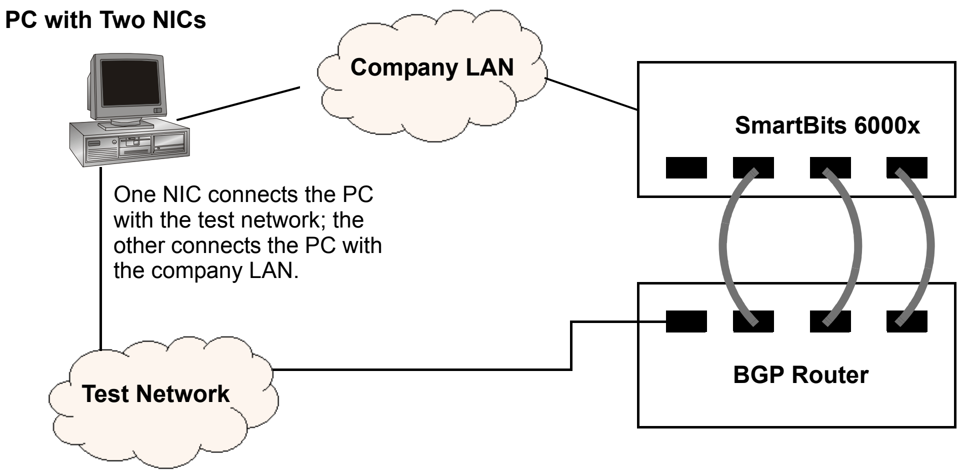

Figure 18-4. PC/Hardware Configuration for a PC with Two NICs

Connecting a PC to the BGP Router

In order for the SmartFlow BGP feature to advertise routes to a router, the PC must be set up to communicate with the router. Note that the PC may be connected to a BGP router either directly or through a test network. Procedures required to configure the router vary according to the type of router used. An example procedure is as follows:

To connect the PC to a router:

1. Enable **bgp** on the router.
2. Enable **bgp neighbor** on the router and indicate the IP address of the neighbor system.
3. Check the router status. From the command line on the router, type a command such as:
   - **Show ip bgp** to show all bgp routes.
   - **Show ip bgp neighbor** to show the status of the neighbor connection.

Note: The router commands may vary from the commands in this procedure.

Setting up BGP Routes

For the SmartFlow BGP feature, the networks set up in SmartFlow are the routes that will be advertised. The networks set up on the *IPv4 Networks* tab are mainly for flow creation, but are also used to advertise routes. For more information, see *"Defining IPv4 Network Information for a Port" on page 106*. To add more networks that are used exclusively to test the capacity of the device and are not related to flows, use the *BGP* tab. Refer to the online Help for more information.

You can set up as many as 100,000 networks per port in SmartFlow. To set up many routes, once the initial networks are established on the *IPv4 Networks* tab, use the *Additional Routes* fields on the *BGP* tab.

Port Setup

When setting up ports, the fields on the *Cards* tab that are unique to the BGP feature are *BGP* and *AS ID*. (See *Figure 18-5*.)

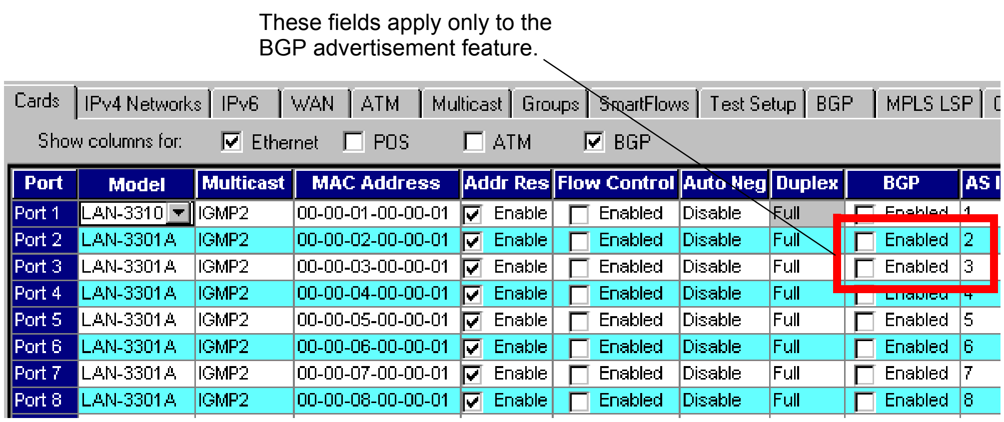

Figure 18-5. Cards Tab with BGP-related Fields

For a description of each field on the *Cards* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Specifying BGP Parameters

The BGP parameters that need to be assigned involve PC host setup, router setup, and route advertisement and removal.

To assign BGP parameters, click the *BGP* tab. (See *Figure 18-6*.)

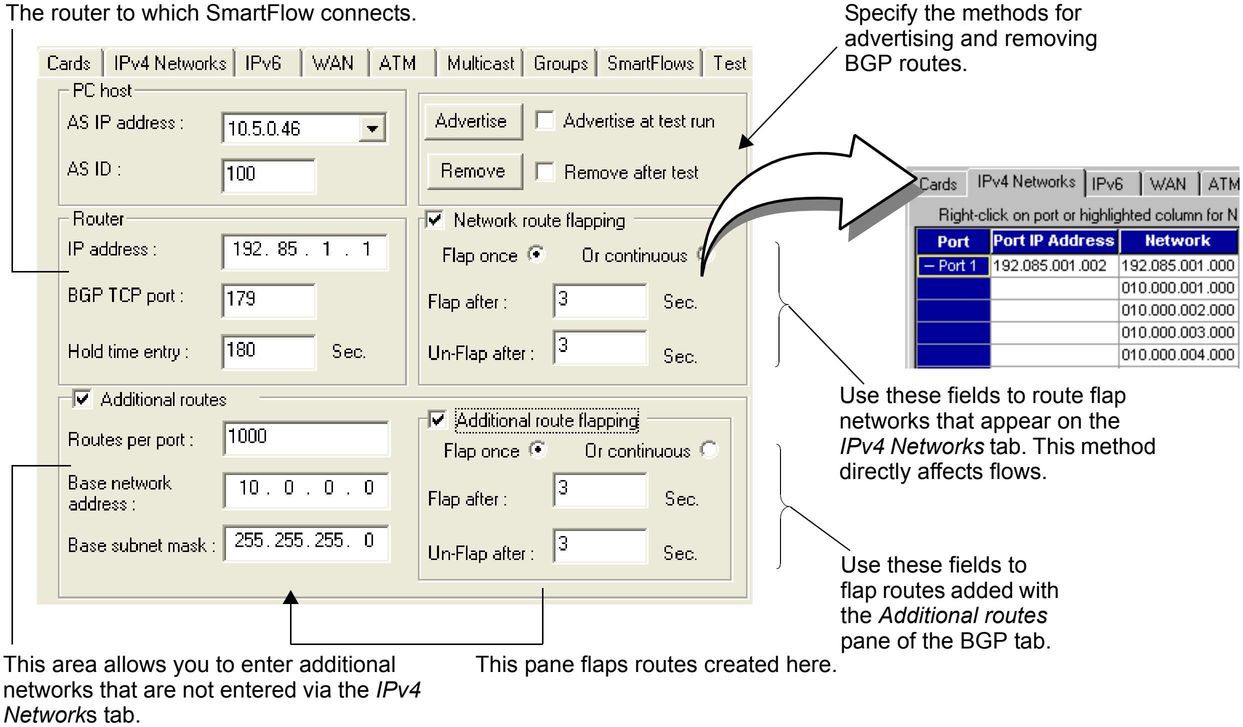

Figure 18-6. BGP Tab

Route Flapping and Unflapping

You can flap routes using one or both of these two methods:

- Network Route Flapping

  This method flaps the networks (routes) set up on the *IPv4 Networks* tab. It *directly* affects flows because flows are based on the network addresses in the *IPv4 Networks* tab. Once network routes are flapped (removed), the router does not know where to forward the flows. This causes the route to drop frames from flows and also increases latency.

- Additional Route Flapping

This method flaps the additional routes set up in the *Additional routes* pane of the *BGP* tab. Since no traffic is sent to these routes, it *indirectly* affects flows by causing the router to work harder and increases latency with some frame loss.

If the *Or continuous* button in the *Additional route flapping* pane is clicked, SmartFlow continues to flap network and/or additional routes until each unflap time has elapsed, at which point SmartFlow re-advertises the routes.

If a test is stopped before it finishes and SmartFlow is in the middle of advertising routes, the test does not stop until it advertises all routes.

Important: The test duration (set on the *Test Setup* tab) should be large enough to accommodate both route flapping and unflapping times.

If you flap routes, you cannot flap them again without unflapping them. Therefore, if the *Or continuous button* was clicked, you should also enter a value in the *Un-flap after* field.

See *"Examples of Flap Settings" on page 458* for a better understanding of how the various settings affect flapping.

These are the parameters on the *BGP* tab in the *Additional route flapping* pane that relate to both route flapping methods: the *Flap once* button, the *Or continuous* button, the *Flap after* field, and the *Un-flap after* field. For a description of each field on the *BGP* tab, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

Notes:

- Flap and unflap times should be long enough to allow a device to stabilize after a flap or unflap. Spirent Communications recommends 30 seconds for each of these values. However, your test environment and testing needs may differ.
- If you schedule flapping during a Latency SnapShot test, Spirent Communications recommends setting at least a 2-second difference between the first flapping and the time to start collecting shapshot data.

Examples of Flap Settings

Here are two examples of flap and un-flap time parameters.

Example 1: Flap once

Let us say you set flap parameters as shown here:

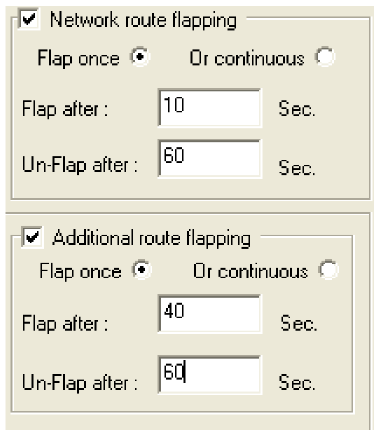

In these example settings, the test duration is 120 seconds. According to these settings, the route flapping is performed as listed here. All times are from the start of the test.

1. 10 seconds after the test starts:
   SmartFlow flaps network routes.
2. 40 seconds after the test starts:
   SmartFlow flaps the additional routes set up.
3. 70 seconds after the test starts:
   SmartFlow *un*flaps network routes.
4. 100 seconds after the test starts:
   SmartFlow *un*flaps additional routes.

Example 2: Flap continuous

Let us say you set flap parameters as shown here:

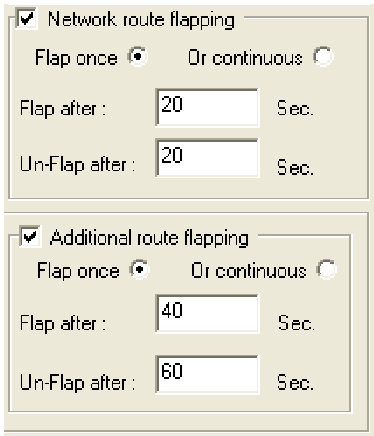

In these example settings, the test duration is 120 seconds. According to these settings, the route flapping is performed as listed here. All times are from the start of the test.

1. 20 seconds after the test starts:
   SmartFlow flaps network routes.
2. 40 seconds after the test starts:
   SmartFlow *un*flaps network routes and flaps the additional routes set up.
3. 60 seconds after the test starts:
   SmartFlow flaps network routes again.
4. 80 seconds after the test starts:
   SmartFlow *un*flaps network routes.
5. 100 seconds after the test starts:
   SmartFlow flaps network routes and *un*flaps additional routes.

Interpreting BGP Test Results

Results for the BGP feature of SmartFlow are just like test results in any SmartFlow test.

When doing BGP testing, look at results for this behavior:

- For tests run without advertising, test results should yield 100% frame loss.
- For tests run with advertising, at least some frames should transport from one route to another.

Chapter 19

IP Multicast Testing

This chapter contains information about the multicast feature of SmartFlow and general information about multicast technology.

In this chapter...

What is IP Multicast?

IP multicast is a method of transmitting traffic (text, audio, and video) on the Internet or an internal network from a sender to multiple receivers. It is a one-to-many transmission method because the server sends a single multicast packet to multiple multicast group recipients simultaneously.

As a networking protocol, IP multicast is a more efficient method for multipoint delivery of IP datagrams than unicast (sending multiple copies of the same frame) and broadcast (sending the frame to all network nodes).

IP Multicast Applications

Typical IP multicast applications include:

- Multimedia, including:
  - Streaming media
  - Web-based training
  - Corporate broadcasts
  - Desktop conferencing
  - Live real-time videocasts over the Internet
  - Live TV and radio to the desktop
- Data warehousing
- One-to-many data PUSH applications
- Real-time news bulletins, weather, and stock quotes
- Database and cache updates
- File transfers from a central distribution database to remote field stations
- Whiteboarding in workgroups.

Any collaborative computing that requires simultaneous communications between groups of computers is suitable for multicast transmission. For example, the London Stock Exchange requires that each host in the group can both transmit and receive information, so that every host has the most updated information.

IP Multicast Basics

Multimedia and data sharing applications require IP multicast functionality in network devices and networks.

About IGMP

Internet Group Management Protocol (IGMP) is the standard protocol for IP multicasting on the Internet. It is used for the initial setup of the multicast session and does not carry data. IGMP is used to set up host memberships in a multicast group on a single network. The host informs its local router that it wants to receive messages with the address for a particular multicast group. Hosts that conform to level 2 of the IP multicasting specification must use IGMP. IGMP is defined in RFCs 1112 (Version 1) and 2236 (Version 2).

How IP Multicast Sessions Work

Local routers periodically multicast IGMP membership queries to 224.0.0.l, the “all hosts” multicast address, to determine group memberships. Only one member per group per subnet responds to the query with a report, while other members suppress reports.

To receive multicast traffic, a host (i.e., a receiving port that can be a host, PC, or workstation) must request the multicast application to join the multicast group at the local router by using IGMP. The application on the host sends an IGMP membership report join request to the local router. The request is processed through the IP stack. The router queries the host for membership report using the destination address of 224.0.0.1, with a Time to Live (TTL) value of 1. Each host in the group responds with a membership report.

A host can join and leave a group at any time. A host can be a member of more than one group.

How IP Multicast Works

Multicast packets only travel along paths that lead to recipients that have requested to be part of the multicast group (transmission), which also cuts down on traffic, as illustrated in *Figure 19-1 on page 464*. It is the technology used for transmitting webcasts such as live, on-demand audio, and video broadcasting.

A multicast host only sends traffic to routers whose endpoints are receivers that requested the multicast transmission, cutting down on traffic.

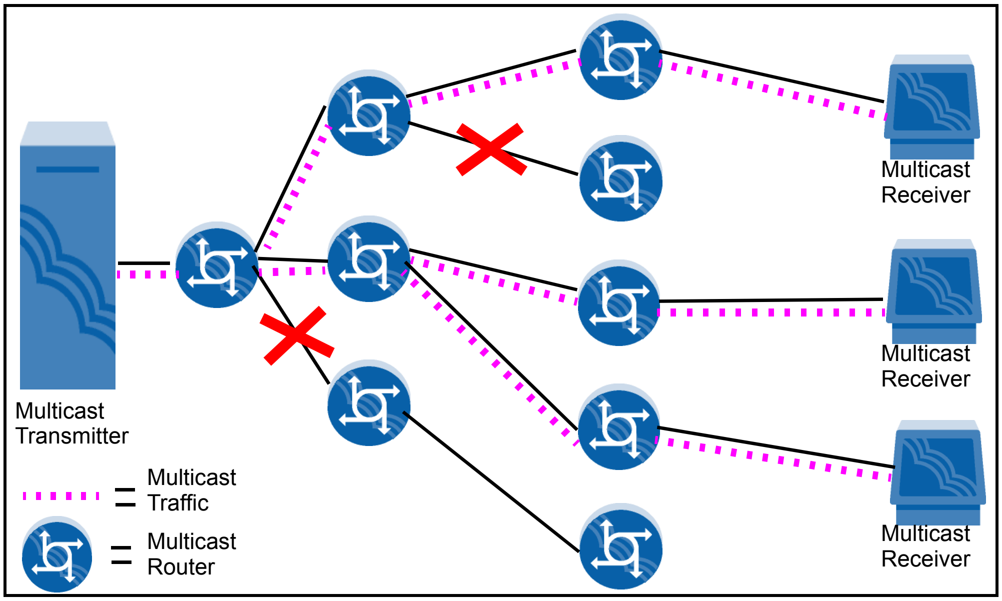

Figure 19-1. Multicast Transmitter Only Sends to Routers with Receivers that Joined Group

IP Multicast Addressing

IP multicast is based on the use of multicast addresses and address groups. These are Class D IP addresses in the range 224.0.0.0 to 239.255.255.255, as assigned by the Internet Assigned Numbers Authority (IANA). *Table 19-1* and *Table 19-2* show the IP multicast address ranges and the registered group addresses.

Table 19-1. Reserved Address Ranges

Address Range	Reserved for...
224.0.0.1 to 224.0.0.255	Routing protocols
239.0.0.0 to 239.255.255.255	Local intranets

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP Multicast groups. *Table 19-2* lists well-known multicast groups.

Table 19-2. Well-known Multicast Groups

IPv4 Address	Multicast Group
224.0.0.1	All systems on this subnet

Table 19-2. Well-known Multicast Groups (continued)

IPv4 Address	Multicast Group
224.0.0.2	All routers on this subnet
224.0.0.4	All DVMRP routers
224.0.0.5	All OSPF routers
224.0.1.11	IETF-1-Audio
224.0.1.12	IETF-1-Video
224.0.1.7	Audio news
224.0.1.16	Music service

How IP Multicast IP Addresses Map to Ethernet Multicast Addresses

IP multicast addresses map to Ethernet multicast addresses in the following ways:

- The lower 23 bits are mapped one-to-one.
- The MAC address prefix is 01-00-5E (a well-known MAC address for IP multicast), plus the last 23 bits of the IP multicast address.

  Example 1: 224.100.50.25 in hex is E0-64-32-19, so the MAC address is 01-00-5e-64-32-19

  Example 2: 224.241.50.25 in hex is E0-F1-32-19, so the MAC address is 01-00-5e-71-32-19 (The 24th bit is not used.)

Note: Several IP multicast addresses can map to the same Ethernet multicast address. For example, each of these IP multicast addresses can map to the same Ethernet address of 01-00-5e-0a-00-01: 224.10.0.1, 225.10.0.1, 226.10.0.1, 239.138.0.1.

Why Test IP Multicast?

The products that support IP multicast vary significantly in functionality and performance. As the use of IP multicast becomes more common and more devices use the Internet Group Management Protocol (IGMP), you need to understand the effect it has on the IP multicast devices and networks. For example, does the device process IP multicast traffic the same way it processes unicast traffic?

Why Use VLANs in Multicast Testing?

Packets that originate from a VLAN contain a VLAN tag. (Routers treat packets from non-VLANs as untagged VLAN packets.) The ability to send tagged and untagged VLAN packets emulates a VLAN trunk connection or many multicast members on a single port/link, each using multiple VLAN IDs. *Figure 19-2* illustrates the latter.

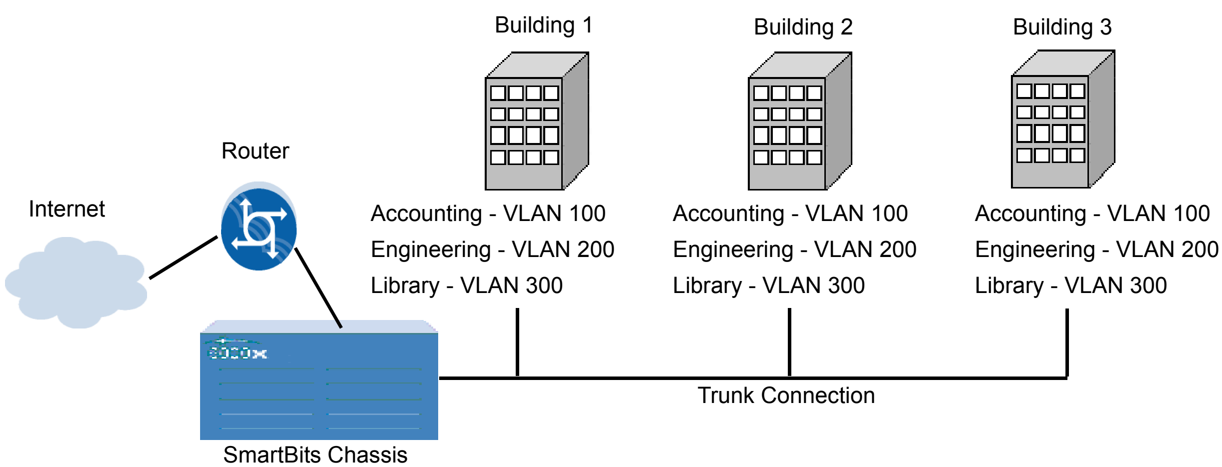

Figure 19-2. Why Include VLAN Tags in Multicast Traffic Testing

In *Figure 19-2*, if the accounting department wants to send traffic to the engineering department, the traffic must go through the router. Only the incoming traffic through the switch for the VLAN has the VLAN tag.

Thus, it is important to be able to test multicast traffic with and without VLAN tags since multicast traffic may originate from a port that is in either a tagged or non-tagged VLAN.

What is the IP Multicast Feature of SmartFlow?

SmartFlow helps verify the IP multicast functionality and features of your IP multicast devices or networks. Traffic types (unicast or multicast) may be configured independently across different streams on the same port. This allows mixed traffic generation.

SmartFlow allows you to test:

- The effect of the IP multicast application on the network by testing mixed-class (unicast and multicast) throughput. Use the same SmartBits port to generate mixed traffic.
- The effect of using multiple groups.
- Latency to determine the ability of a port to handle voice or video.
- The effect of the IP multicast application on the network by testing mixed-class throughput.
- The effect of using multiple groups.
- Latency to determine the ability of a port to handle voice or video.
- IGMP Versions 1 and 2.

SmartFlow supports only IPv4 for multicast traffic. It supports both IPv6 and IPv4 only for unicast traffic.

The maximum number of multicast group/VLAN combinations allowed per port is 1,024.

Notes:

- You can use cyclic flows, but only for the transmitter of the multicast group (by varying the *SRC IP address*, *SRC MAC address*, or *SRC port* fields).
- SmartFlow sends ARP requests for unicast flows only.
  Use the *Test Setup>Sample Iteration* tab options for multicast tests.

Card/Module Support for Multicast Traffic

Any card/module that SmartFlow supports can act as a multicast group transmitter. Multicast group receiver ports must be capable of sending IGMP messages (join or leave requests). The following cards can send IGMP messages and thus act as multicast receivers, with both tagged and untagged VLAN support:

- ML-5710A
- ML-7710/7711
- LAN-3101A/B, LAN-3102A, LAN-3111A
- LAN-3201x
- LAN-33xx
- LAN-3306 XD and LAN-332x XD
- XLW-3720A/XLW-3721A

- XFP-3730A/XFP-3731A.

IP Multicast Maximums in SmartFlow

SmartFlow supports IP multicast as follows:

- Up to 7,000 (LAN-3201x) multicast groups per transmitter.
- Depending on card type, up to 2,047 multicast group/VLAN combinations per port. This number is based on the ability of the card to send join/leave requests. A multicast group/VLAN combination is a multicast group plus untagged or tagged VLANs. For example, if there are five multicast groups, each with 20 VLANs, there are 100 combinations. See *Table 19-3 on page 468* for more information.
- In cyclic flows, up to 65,536 variations of each flow's source IP address *only.* This is the same value as for unicast traffic, but in multicast traffic it simulates numerous transmitters. The multicast group address and MAC address remain constant.

Table 19-3. Maximum Multicast Group/VLAN ID Combinations by Module Type

Module	Maximum Number of IGMP Group/VLAN ID Combinations per Receiving Port	Maximum Number of IGMP Groups per Transmitting Port
LAN-3101A/B[1] LAN-3102A LAN-3111A ML-7710 ML-7711 ML-5710A	999	999
LAN-3201B LAN-3201C	1002	7000
LAN-3300A LAN-3301A LAN-3302A LAN-3310A LAN-3311A XFP-3730A XFP-3731A XLW-3720A XLW-3721A	511	511

Table 19-3. Maximum Multicast Group/VLAN ID Combinations by Module Type

Module	Maximum Number of IGMP Group/VLAN ID Combinations per Receiving Port	Maximum Number of IGMP Groups per Transmitting Port
LAN-3306A LAN-3320A LAN-3321A LAN-3324A LAN-3325A LAN-3327A	2047	2047

1 If jumbo/oversized (up to 16,388 bytes) frames are enabled, the maximum number of multicast groups for the LAN-3101A/B module (whether it is a transmitting or receiving port) is reduced to 300.

Note: When using cyclic flows to increase the number of IGMP groups per transmitting port, the maximum number of groups allowed increases to the cyclic flow maximums. See *Table 2-6 on page 45* for these cyclic flow maximums by card type.

How IP Multicast Works in SmartFlow

In SmartFlow, the DUT makes a copy of the packets in the SmartFlow and sends them to each receiving multicast port. Although the DUT copies the packets, resulting in multiple streams, only one trackable UDP SmartFlow exists for a multicast group in which only one transmitter exists. This is shown in *Figure 19-3*. If the multicast group has multiple transmitters (e.g., the London Stock exchange in which each receiver can also transmit), there are as many SmartFlows as there are transmitters in a multicast group.

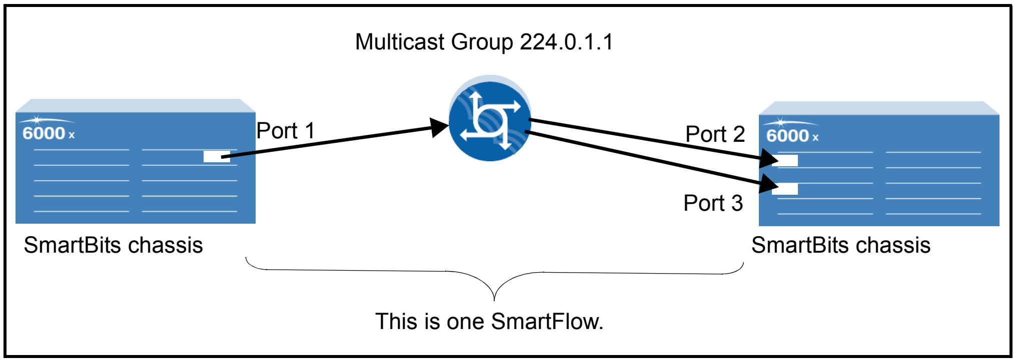

Figure 19-3. Multicast Traffic Only Uses One SmartFlow per Transmitter in a Group

The SmartBits card on which the transmitter port resides generates traffic with a separate stream for each receiver port in a multicast group.

Note: IGMP join and leave requests are not counted in SmartFlow results.

Think of the multicast group IP address as the destination address for the SmartFlow. Once multicast flows are created, the *SmartFlows* tab shows the multicast group IP address instead of the destination port for each multicast SmartFlow. This is shown in *Figure 19-4*.

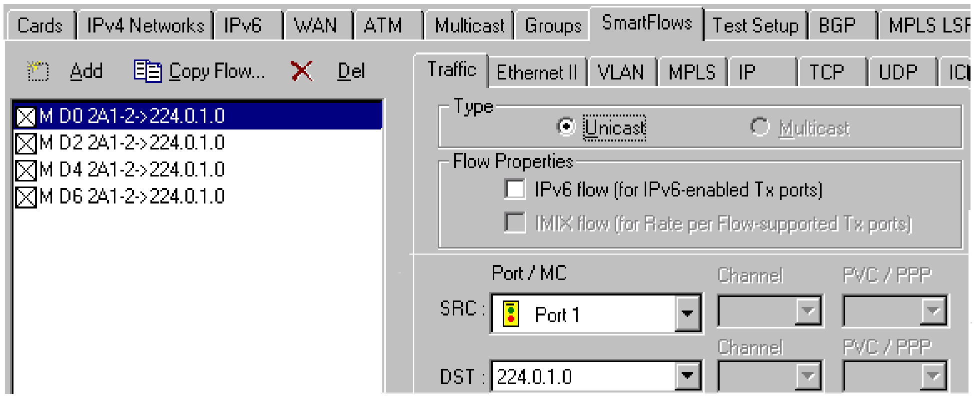

Figure 19-4. Multicast SmartFlows show the group address as the destination port

Multicast Test Phases

When a test includes multicast traffic, the test runs through much of the same phases as a test with unicast traffic. For a description of the test phases, see *"Test Phases" on page 50*.

When you run a test containing *only* multicast traffic, these differences exist:

- No learning phase takes place.
- Receivers of the multicast group send out join requests just prior to the start of the test.
- If you selected to send leave requests at the end of the test, the receivers perform this task at the end of the test.

When you run a test containing multicast *and* unicast traffic, these differences exist:

- Learning takes place for the unicast traffic.
- Receivers of the multicast group send out join requests just prior to the start of the test.
- If you selected to send leave requests at the end of the test, the multicast group receivers perform this task at the end of the test.

How to Test IP Multicast

When testing multicast, you may want some of these questions answered:

- What is the effect of the new multicast application on day-to-day critical network?
- At what load can the network handle IP multicast?
- What is the effect of multicast traffic on other applications?
- How many types of IP multicast applications (groups) can the network handle concurrently?
- Is different/more equipment required?

SmartFlow can help answer these questions by setting up tests according to any of these scenarios, which represent typical network/device conditions for multicast. Listed by increasing complexity, these tests are:

- *Multicast only traffic (no unicast), with and without VLAN tagged traffic.*
  Use this scenario just to test the ability of a device to handle multicast traffic only.
- *Mixed unicast and multicast traffic.*
  Use this scenario just to test the ability of a device to handle multicast traffic in the presence of unicast traffic.
- *Mixed transmitters and receivers across multicast groups.*
  An example of this test is the London Stock Exchange, where each port in the group is both a transmitter and a receiver.
- *Mixed unicast and multicast traffic plus mixed transmitters and receivers.*
  The unicast traffic would go to both the transmitting and receiving ports of the multicast group.

General Test Setup for IP Multicast Testing

Setting up a test for multicast traffic includes all of the steps performed when setting up any other test. For general information about how to set up a test, see *Chapter 8, "Set up and Run Tests."*

These are additional steps relating specifically to multicast testing that must be done:

1. Make sure that all router ports to be used in the test are configured for the correct version of IGMP.
2. At the *Cards* tab, specify the version of IGMP to be used in the *Multicast* field. (IGMP versions can differ across ports.) The DUT must also support the version that is selected.

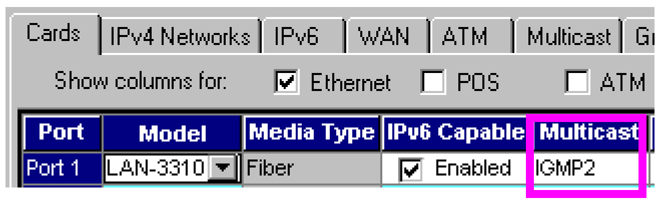

   See *"Card/Module Support for Multicast Traffic" on page 467* for list of cards that support IGMP.
3. If you want the traffic transmitted to contain VLAN tags, on the *IPv4 Networks* tab, add any VLANs to the network(s) that you plan on using for the multicast group transmitter(s).
4. At the *Multicast* tab, define the multicast group IP addresses and any needed for VLANs for multicast group receiver ports.

   (The transmitter of the multicast group is defined when the flows are set up.)
   For more information, see *"Adding Multicast Groups and VLANs" on page 477*.
5. Use the *Group Wizard* to set up multicast flows, or manually set them up using the *SmartFlows>Traffic* tab.
   Make sure to select multicast as the traffic type by clicking the *Multicast* button. The transmitting port of the flow determines the transmitter of the multicast group. For more information, see *"Creating Multicast Flows and Groups" on page 487*.
6. On the *Test Setup>Sample Iteration* tab, specify the option to run a sample iteration with a confirmation. Do this if you want to test the validity of your configuration (without running the entire test), ensure that multicast groups were successfully set up, and verify multicast route entry-related aspects of the DUT (such as Layer 3 outgoing interface lists).
   For more information about this tab, see *"Sample Iteration Tab" on page 222*.

7 On the *Global Multicast Properties* dialog box, enter a delay in the **Delay after IGMP Join requests (Sec.)** field that is large enough to allow the DUT to process the join requests from all receiving ports. You can also specify a delay between join or leave requests.

For more information about this dialog box, see *"Controlling IGMP Join/Leave Requests and Traffic Timing" on page 483*.

8 Run the test and view results.

Defining Multicast Groups and VLANs

From the *Multicast* tab (*Figure 19-5*) you can define most of the elements of the multicast group, such the multicast group IP address and multicast group receiver ports. Define the multicast groups and VLANs at this tab before creating multicast flows and SmartFlow groups.

The *Multicast* tab displays:

- IP address(es) of multicast group(s).
- Receiver ports that support multicast traffic and on which multicast is enabled (on the *Cards* tab).
- Which receiver ports join a multicast group (and send join/leave requests, if selected).
- Any VLANs associated with receiver ports and whether the VLANs are part of the multicast group (and send join/leave requests, if selected).

Note: If no ports support multicast or a version of IGMP for the port at the *Cards* tab was not selected, no ports are displayed on the *Multicast* tab.

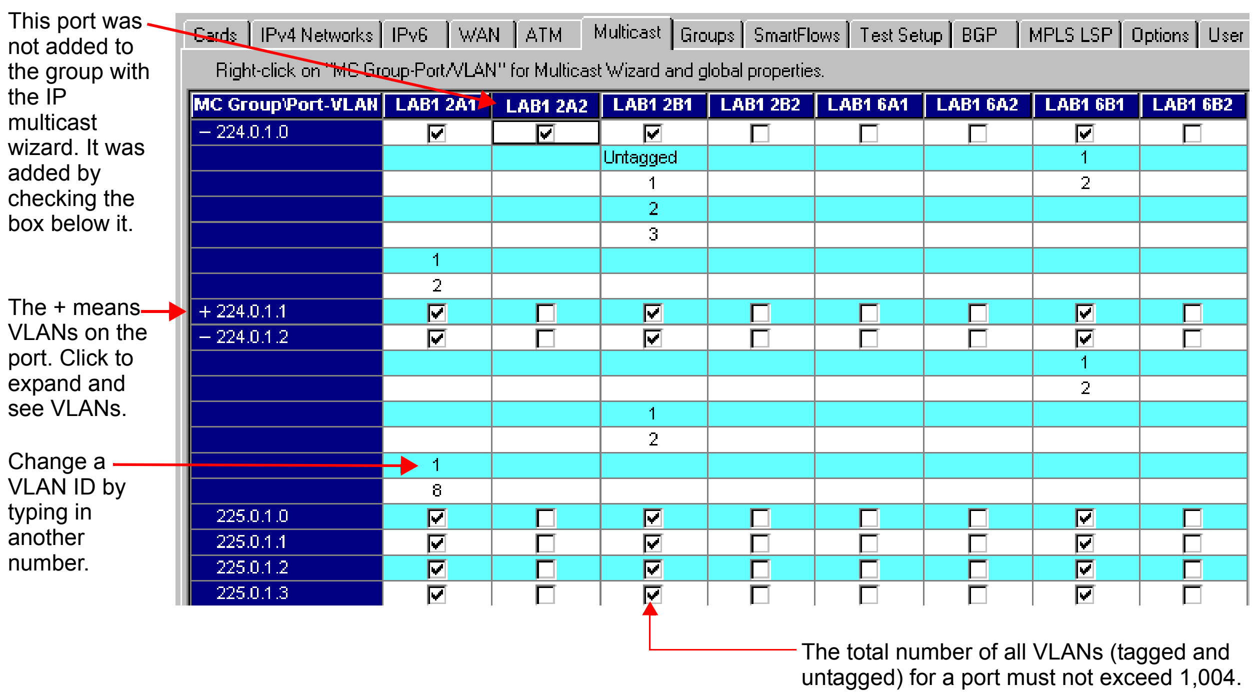

Figure 19-5. Multicast Tab with Multicast Group Addresses and VLANs Displayed

To make a port part of a multicast group, click the corresponding port box that is located in the same row as the group.

If a multicast group has VLANs associated with it, a + sign appears to the left of the group IP address. To view all VLANs associated with a port, click the + sign.

Untagged indicates a VLAN with no tag. A checked port with no VLAN IDs has an assumed untagged VLAN. Once you add VLAN IDs to a port with an untagged VLAN, the word “Untagged” displays.

To expand or collapse all VLANs of all multicast groups, right-click in the far left column. Then select **Expand All** or **Collapse All** from the pop-up menu.

Notes:

- Depending on card type, the maximum number of multicast groups/VLAN combinations allowed per receiving port is 2,047. See *Table 19-3 on page 468*.
- On LAN-3201 cards, each port supports only one tagged VLAN or an untagged VLAN per multicast group.
- ML-5710 and ML-7710/11 cards only support one VLAN tag per multicast group on the receiver ports.

If you defined ports as multicast ports (on the *Cards* tab) but did not yet define any multicast groups, the tab only displays the multicast-enabled ports. *Figure 19-6* shows what the *Multicast* tab might look like:

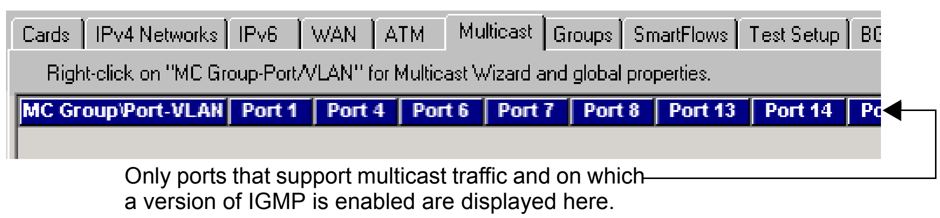

Figure 19-6. Multicast Tab Prior to Defining Any Multicast Groups

Defining Multicast Group Transmitters

Since the multicast group transmitter is actually the transmitting port in each of the flows, the transmitter is defined when the flows are set up (using either the Group Wizard or the *SmartFlows* tab). To define VLANs for the transmitter of the group, use the *IPv4 Networks* tab. (See *“Defining IPv4 Network Information for a Port” on page 106* for more information.)

Transmitting ports can support up to 7,000 (LAN-3201) multicast groups. For cyclic flows, this number can be much higher.

Tip: To have a transmitter port also receive traffic from itself, select the port on the *Multicast* tab (or add an untagged VLAN if there are VLANs on the port) so that it is also a multicast group receiver port. Make sure to do this before flows are set up, since the Group Wizard does not allow you to specify this.

Adding Multicast Groups and VLANs

When testing using multicast traffic, you probably want many multicast groups each consisting of many receiver ports. Each VLAN that is part of the multicast group sends join requests to the transmitter along with the base receiver ports of the group.

A receiver port can join up to 1,024 multicast groups. See *"IP Multicast Maximums in SmartFlow" on page 468* for more information on the maximum number of groups allowed per transmitter port.

Depending on card type, you can add up to 2,047 multicast group/VLAN ID combinations. (See *Table 19-3 on page 468*.) Each tagged or untagged VLAN ID counts as a receiver. One multicast group with one VLAN ID is one combination. For example, if you have 100 VLANs on a port and each VLAN joins every multicast group, there are only have 10 multicast groups.

Notes:

- To make a port the transmitter of a multicast group, define it as the transmitting port when you generate flows (using either the Group Wizard or the SmartFlows tab).
- Any VLAN tags that you add to ports on the *Multicast* tab are included only in the IGMP join and leave requests sent by the port, not in the normal test traffic. To add VLAN tags to normal test traffic, add VLANs to ports at the *IPv4 Networks* or *IPv6* tab.

The IP Multicast Wizard allows you to quickly add one or multiple IP multicast groups, and any VLANs you want to associate with the multicast group's receiver ports.

Tip: If you want to quickly create a new multicast group from an existing group with the same ports, right-click on the specific group address in the **MC Group\Port-VLAN** column. The ports are already selected when the IP Multicast Wizard opens. Make any selections as needed.

For a detailed procedure on how to add multicast groups and/or VLANs, refer to the online Help.

Using the IP Multicast Wizard

Use the IP Multicast Wizard to populate the *Multicast* tab with multicast groups and any VLANs that you want to associate with receiver ports. *Figure 19-7* illustrates this scenario:

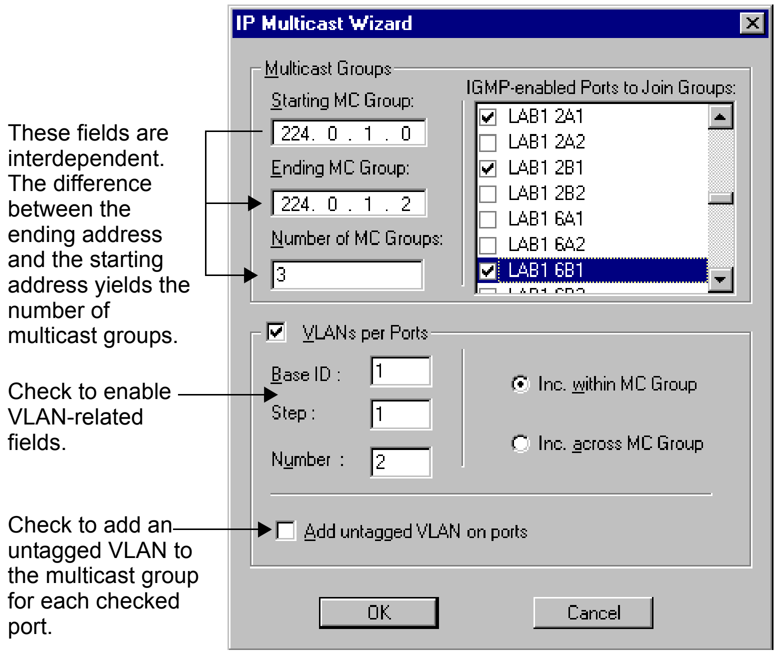

Figure 19-7. IP Multicast Wizard

Important: Once multicast groups are created using the IP Multicast Wizard, if the wizard is used subsequent times to add more VLANs and/or groups and the new range includes existing group addresses, the existing port/VLAN configuration for the groups is overwritten by the new configuration. To add more VLANs to a group and maintain the existing VLANs, use the *VLAN for Multicast Group* dialog box instead. See *"Adding VLANs to a Single Multicast Group" on page 480* for more information.

For a description of how to access the IP Multicast Wizard as well as a description of each field in the *IP Multicast Wizard* dialog box, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

To add more ports to the group, check the box under the port number on the *Multicast* tab.

Example:
According to the setup shown in *Figure 19-7*, three multicast groups will be set up, each consisting of the three receiver ports and two VLANs per port. This means that each port will send six unique IGMP join requests to join the multicast group when the test is started.

Figure 19-8 shows the resulting multicast groups and VLANs that are created and displayed on the *Multicast* tab (based on the setup shown in *Figure 19-7 on page 478*).

Ports selected when the multicast groups were added with the IP Multicast Wizard are already checked.

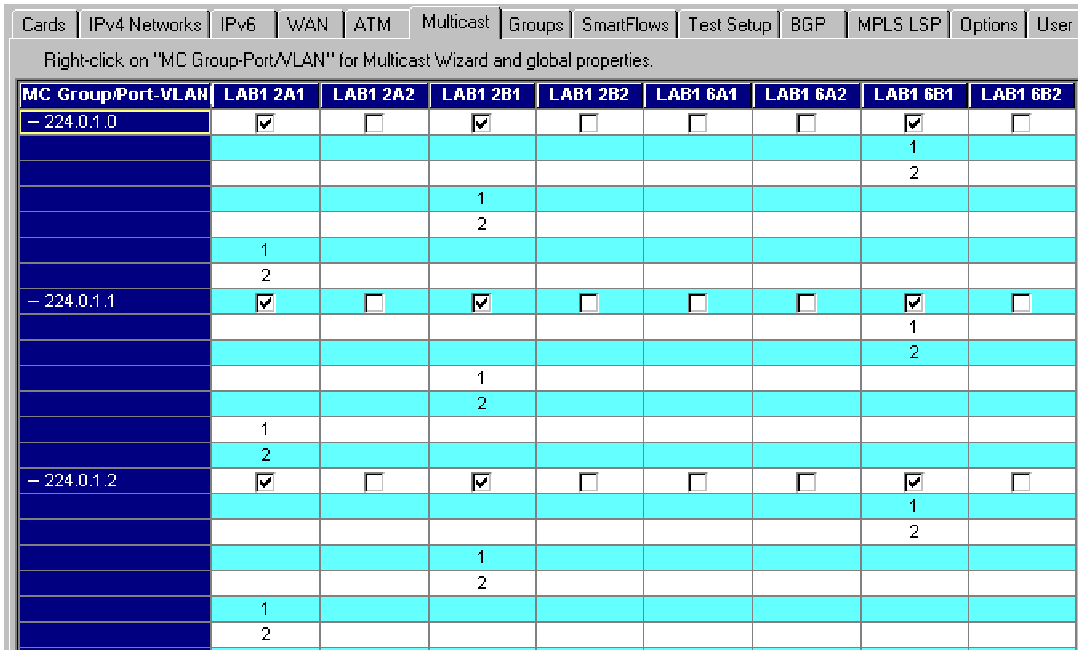

Figure 19-8. Example of Multicast Tab

How MC Group IP Address Fields Work

When you add more than one multicast group using the IP Multicast Wizard, SmartFlow first increments the least significant (right-most) bit in the starting IP address, then the next bit, and so on as needed.

Specify the number of multicast groups to add via the IP Multicast Wizard in any of these ways:

- Enter a value greater than 1 in the **Number of MC Groups** field.

  The address in the *Ending MC Group* field automatically increases so that it reflects the specified number of groups.

- Increase the address in the **Ending MC Group** field.

  The value in the *Number of MC Groups* field automatically increases so that it reflects the resulting number of groups.

- Increase the address in the **Starting MC Group** field.

  The value in the *Ending MC Group* field increases by the number of groups specified in the *Ending MC Group* field.

Note: If the maximum number of addresses already exists, SmartFlow tries to lower the address in the *Starting MC Group* field. If this is not possible, it subtracts the value in the *Starting MC Group* field from the value in the *Ending MC Group* field.

Adding VLANs to a Single Multicast Group

Once a multicast group is set up, VLANs can be added to it. Even if you created the group and added VLANs using the IP Multicast Wizard, you can add more VLANs to the group using the *VLAN for Multicast Group* dialog box.

The *VLAN for Multicast Group* dialog box, as shown in *Figure 19-9*, allows you to *append* VLANs to a multicast group with or without existing VLANs.

Note: Verify that you right-clicked on the correct multicast group to which you want to add VLANs.

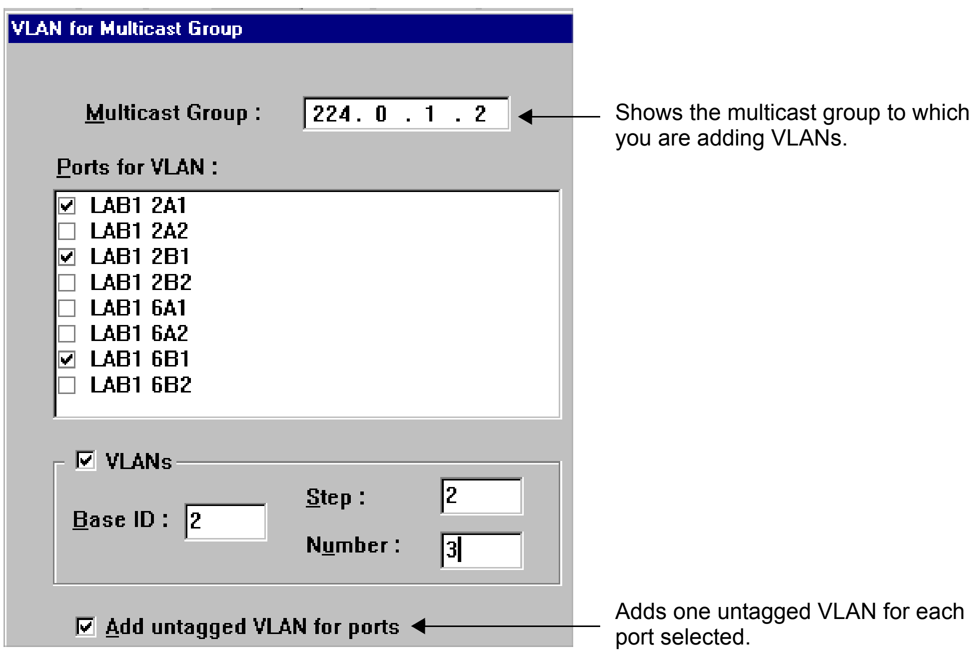

Figure 19-9. VLAN for Multicast Group Dialog Box to Add VLANs to Single Group

For detailed procedures on how to add VLANs to a single multicast group, refer to the online Help. Press **F1** at the tab or select **Help** from the menu bar.

Copying VLAN IDs from One Multicast Group to Another

You can copy rows of VLAN IDs on the *Multicast* tab from one multicast group and paste them to another. When you add multiple VLANs at one time using the *IP Multicast Wizard* or the *VLAN for Multicast Group* dialog boxes, the VLANs are consecutively numbered or increment in regular steps. However if you manually add or change VLAN IDs so the IDs are irregularly numbered (such as VLANs 2, 8, and 57), the ability to copy and paste them to another multicast group is particularly useful.

Important: If the number of rows that you are copying *from* is greater than the number of existing rows you are pasting *to*, VLAN IDs for some ports may be overwritten. It is best to copy all of the VLANs for all ports in a multicast group to another multicast group.

Refer to the online Help for a procedure to copy the VLANs from one group to another.

Including and Excluding Receiver Ports/VLANs in Multicast Groups

Once you have created multicast groups and added any VLANs IDs to receiver ports at the *Multicast* tab, you may want to modify which ports and/or VLANs are part of the multicast groups.

Including Receiver Ports/VLANs

To include a receiver port with no VLAN IDs (untagged VLAN) in a multicast group, click the **Multicast** tab. At the *Multicast* tab, select the checkbox that corresponds to the port and multicast group to be included. (See *Figure 19-10*.) If any VLAN IDs already exist for the port, delete these. See *"Deleting Multicast Groups and VLANs" on page 482* for more information.

In *Figure 19-10*, Port 2A1 sends one join request without a VLAN tag (untagged), and three requests with VLAN tags (1, 2, and 4) to join group 224.01.1. Port 2A2 only sends an untagged request to join group 224.0.1.0.

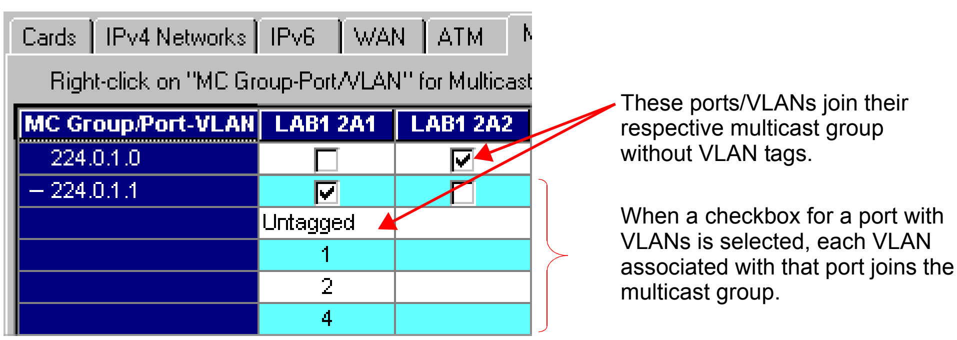

Figure 19-10. Selecting and Deselecting Ports and VLANs

To include all associated VLANs (including an untagged VLAN) of a receiver port, select the checkbox that corresponds to that port and multicast group. To include only tagged VLANs of a receiver port, delete the untagged VLAN.

Excluding Receiver Ports/VLANs

To exclude a port that has no associated tagged VLANs from a multicast group, clear the checkbox that corresponds to that port and multicast group. The port is excluded from joining the multicast group but is not deleted from the *Multicast* tab.

To exclude a port and its associated tagged VLANs from a multicast group, clear the checkbox that corresponds to that port/VLAN IDs and multicast group. A confirmation message displays asking whether you want to permanently delete the VLANs. If you click **Yes**, the port is excluded and all of the VLAN IDs are permanently deleted from the multicast configuration. If you click **No**, the port and its associated tagged VLANs are excluded but not deleted.

Deleting Multicast Groups and VLANs

When you delete a multicast group, any VLANs associated with receiver ports in that group are also deleted. You can delete VLANs independent of the multicast group, and across groups. *Figure 19-11* shows deleting a single VLAN from a receiver port.

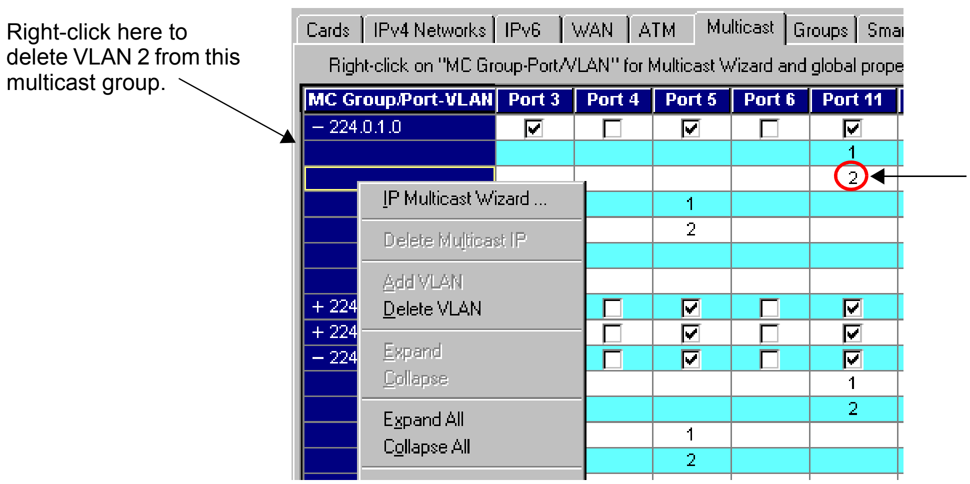

Figure 19-11. Deleting a VLAN from a Receiver Port

For procedures to delete a single tagged or untagged VLAN or multiple VLANs from groups, refer to the online Help.

Setting Up Tests for Multicast Traffic

Tests with multicast traffic are set up in the same way as tests with unicast traffic. However, these differences exist:

- A sample iteration should be used to "train" a DUT for multicast traffic, since learning or ARP requests are not used for multicast traffic. You can also verify group membership using the sample iteration.
  Specify this information on the *Test Setup>Sample Iteration* tab.
- Delays for join requests are available.
  Specify this information on the *Global Multicast Properties* dialog box.
- IGMP leave requests can be sent at the end of the test.
  Specify this information on the *Global Multicast Properties* dialog box.

This section explains the options available for tests with multicast traffic.

Controlling IGMP Join/Leave Requests and Traffic Timing

The Global Multicast Properties dialog box (*Figure 19-12*) allows you to set delays for when IGMP join and leave requests are sent, and if leave requests are sent once the test ends.

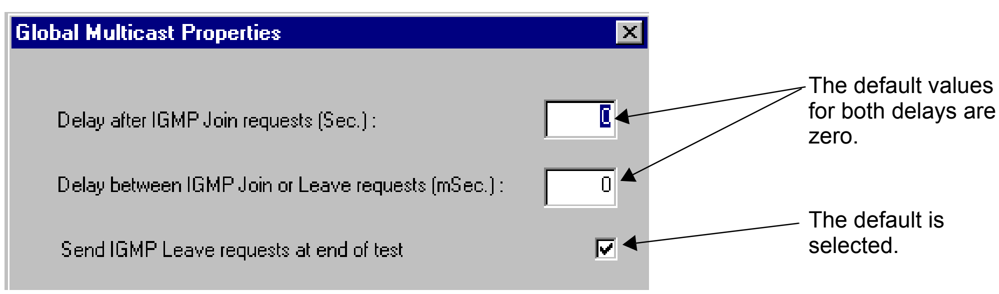

Figure 19-12. Global Multicast Properties Dialog Box

Note: If you want to set a delay for a test with only unicast traffic, use the *Delay before transmit (Sec.)* field on the *Options* tab. If the test contains mixed traffic, the *Delay before transmit (Sec.)* field delays the transmission of both unicast and multicast traffic.

For a description of how to access the *Global Multicast Properties* dialog box as well as a description of each field in the dialog box, refer to the SmartFlow online Help. Press **F1** at the tab or select **Help** from the menu bar.

For additional information about the *Delay between IGMP Join or Leave requests (mSec.)* field, see *"How the Delay between IGMP Join or Leave requests Field Works" on page 484*.

How the Delay between IGMP Join or Leave requests Field Works

How the *Delay between IGMP Join or Leave requests (mSec.)* field works is best understood by an example. Assume you have two receiver ports. Port 1 is a receiver in multicast groups A and C. Port 2 is a receiver in multicast groups B and D. Ports 1 and 2 send out join requests as shown in *Table 19-4*.

Table 19-4. How the Delay Between IGMP Join or Leave requests Field Works

Time Sequence	Receiver Port 1	Receiver Port 2
1	MC Group A - Join Request	MC Group B - Join Request
2	specified delay	specified delay
3	VLAN 1 to MC Group C - Join Request	MC Group D - Join Request
4	specified delay	specified delay
5	VLAN 2 to MC Group C - Join Request	

Notice that in the 5th time sequence, Port 2 did not send any join requests because it already completed sending requests for the groups in which it was a member.

How Delays Affect Mixed Traffic

If a test contains both multicast and unicast traffic, the amount of delay before sending traffic depends on whether or not you specified a non-zero value in each of these fields:

- *Delay after IGMP Join requests (Sec.)* field on the *Global Multicast Properties* dialog box
- *Delay before transmit (Sec.)* field on the *Options* tab.

In mixed traffic tests, these two delays occur back-to-back and affect both unicast and multicast traffic.

After SmartFlow sends out IGMP join requests, a non-zero value only in the *Delay after IGMP Join requests (Sec.)* field causes SmartFlow to delay sending unicast as well as multicast traffic, even though IGMP join requests apply only to multicast traffic. A non-zero value only in the *Delay before transmit (Sec.)* field causes both the unicast and multicast traffic to be delayed by that amount. SmartFlow then sends all traffic simultaneously (rather than staggering the multicast and unicast traffic).

Example:
Assume you have these delays set for a mixed traffic test:
Delay after IGMP Join requests (Sec.): **10**
Delay before transmit (Sec.) on the *Options* tab: **5**

The following sequence of events occurs:

1. Multicast receiver ports send IGMP join requests.
2. Delay for 10 seconds after joins.
3. Delay for 5 seconds before transmitting test traffic.
4. Transmit multicast and unicast test traffic.

Thus, the total delay is 15 seconds before transmitting all traffic.

Delay Lengths for Traffic

To understand exactly how delays set in the *Delay after IGMP Join requests (Sec.)* and *Delay before transmit (Sec.)* fields can affect traffic, see *Table 19-5*. For purposes of example, assume that non-zero values (**10** and **5**) shown in the table are used for delays.

Table 19-5. How Delays Affect Multicast Traffic and Unicast Traffic

Types of Traffic in Test	Delay After IGMP Join (Global Multicast Properties dialog box)	Delay Before Transmit (Options tab)	Total Delay (in seconds)
Multicast, Unicast	10	5	15
Multicast, Unicast	10	—	10
Multicast, Unicast	—	5	5
Multicast	10	5	15
Multicast	10	—	10
Multicast	—	5	5
Unicast	10	5	5
Unicast	10	—	None
Unicast	—	5	5

— Indicates that the value for the delay was left at the default of 0 (zero).

If a test contains at least one multicast flow, the total delay for both unicast and multicast traffic is the sum of the values in the *Delay after IGMP Join requests (Sec.)* field and *Delay before transmit (Sec.)* field. If no multicast flows exist, only the delay in the *Delay before transmit (Sec.)* field occurs.

Training the DUT for Multicast Traffic

Since multicast traffic does not have a learning or ARP phase, use the *Run sample iteration/verify IGMP Joins for Multicast testing* option (on the *Sample Iteration* tab) to allow the DUT to associate the source IP addresses and multicast groups before the first test iteration.

To set access this option, select the **Test Setup>Sample Iteration** tab.

For more information about the *Sample Iteration* tab, see *"Sample Iteration Tab" on page 222*.

Verifying Group Membership

Verify group membership by using the *Run sample iteration/verify IGMP Joins for Multicast testing* option together with the *Confirm packet receipt in sample iteration?* option when the test is set up. When selected, SmartFlow confirms receipt of at least one packet at each receiving port and for each flow destined for the port for multicast traffic.

To set access this option, select the **Test Setup>Sample Iteration** tab.

For more information about the *Sample Iteration* tab, see *"Sample Iteration Tab" on page 222*.

Creating Multicast Flows and Groups

You can set up multicast flows individually using the *SmartFlows* tab, or use the Group Wizard to rapidly set up multiple flows at once. This section explains the Group Wizard traffic patterns available for multicast traffic and their unique behaviors for multicast. See *Chapter 7, "Work with Individual Flows"* and *Chapter 6, "Use the Wizard to Set up Flows/Groups"*, respectively, for more information.

You must define the multicast groups and VLANs on the *Multicast* tab before creating multicast flows and SmartFlow groups.

Note: Do not confuse a SmartFlow group of flows (used for tracking purposes in results) with a multicast group. Whenever possible in this user guide, the word *multicast* precedes the word group when this is the type of group that is intended.

Guidelines for Setting Up Multicast Flows and Groups

These guidelines can help you track multicast flows in test results. It is not necessarily meant that all of these guidelines are used at one time.

- *To track results by multicast group:*
  Create a separate SmartFlow group with a unique name that corresponds to each multicast group.
- *To track results by transmitter*:
  Set a group for each transmitter-to-multicast group combination in a separate flow.
- *To test QoS (or another aspect of traffic) together with multicast*:
  Create a separate SmartFlow group with a unique name for each QoS level. Each of the flows in the SmartFlow group transmits to a different multicast group. To test mixed traffic, create another group for the unicast flows (such as unicast group).
  Example:
  To test three QoS levels with multicast, you use two transmitters, three multicast groups, and the *Backbone* pattern to generate the flows/groups.
  The following table shows the basic setup needed to create this scenario.

SmartFlow Group	Flows in
QoS_0 (contains all flows with QoS 0)	Tx1 > MC Group 1 Tx1 > MC Group 2 Tx1 > MC Group 3 Tx2 > MC Group 1 Tx2 > MC Group 2 Tx2 > MC Group 3

SmartFlow Group	Flows in
QoS_3 (contains all flows with QoS 3)	Tx1 > MC Group 1 Tx1 > MC Group 2 Tx1 > MC Group 3 Tx2 > MC Group 1 Tx2 > MC Group 2 Tx2 > MC Group 3
QoS_7 (contains all flows with QoS 7)	Tx1 > MC Group 1 Tx1 > MC Group 2 Tx1 > MC Group 3 Tx2 > MC Group 1 Tx2 > MC Group 2 Tx2 > MC Group 3

Figure 19-13 shows an example of the resulting groups and flows.

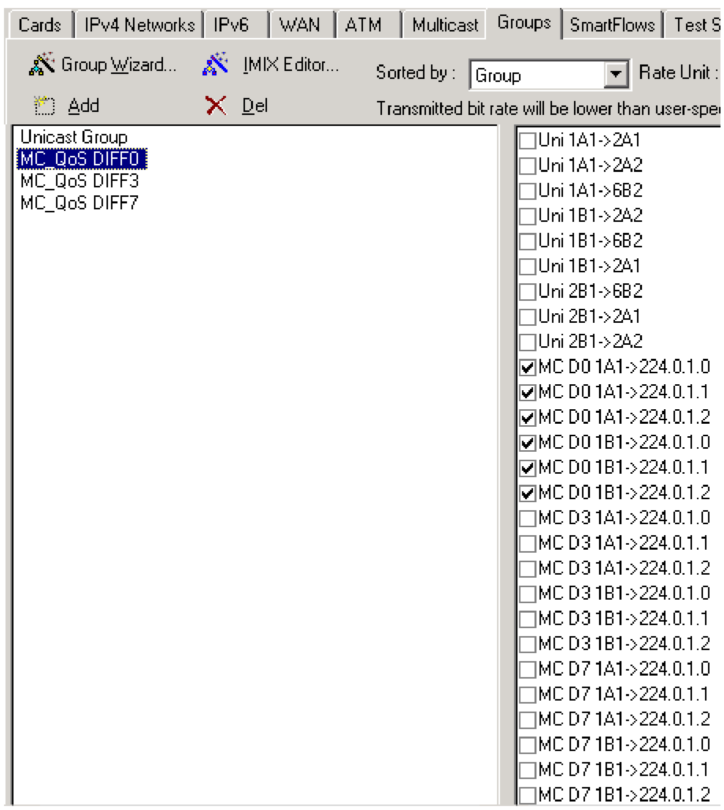

Figure 19-13. Example of Groups for Multicast with QoS and Unicast

Group Wizard Traffic Patterns for Multicast

Use the Group Wizard to create multicast flows and groups the same way you do for unicast traffic. However, once you select the *Multicast* option for the traffic type on the *Traffic Configuration* page of the Group Wizard, the patterns differ from unicast traffic.

Note: The *Waterfall* pattern is not available for multicast traffic.

For purposes of demonstration, there are three multicast groups set up on the *Multicast* tab, each with one untagged and tagged VLANs, and no flows created yet. This scenario is shown in *Figure 19-14*. Each type of traffic pattern example described in this section uses these multicast groups and selected receiver ports that appear on the *Multicast* tab.

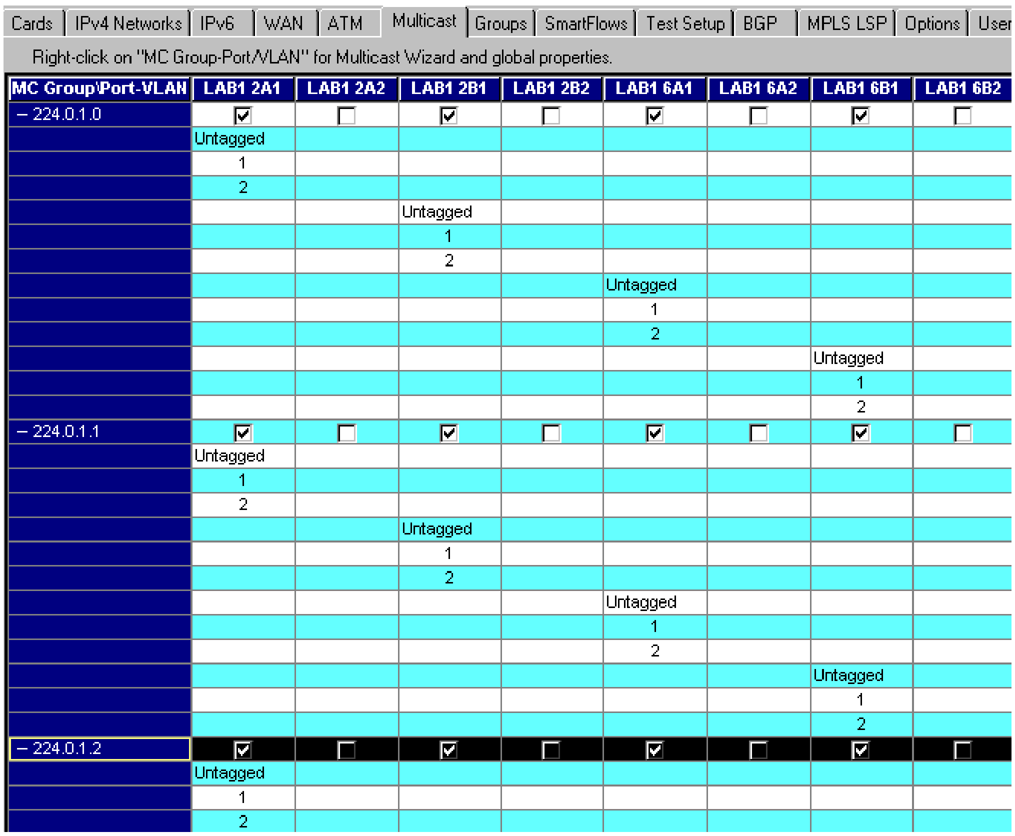

Figure 19-14. Sample Multicast Setup - 3 Groups Each with Tagged and Untagged VLANs

Comparison of Group Wizard Traffic Patterns for Multicast

For multicast traffic, the number of flows that results depends largely on the type of traffic pattern chosen in combination with the number of transmitter and receiver ports.

You can use one pattern to create a number of SmartFlow groups and flows, and then use a different pattern to create more groups and flows. When deciding which pattern(s) to use, keep in mind this information:

- The *Backbone* pattern allows the same transmitter to send traffic to multiple multicast groups. This pattern results in as many flows as there are unique transmitter-to-group combinations.

  This pattern produces the most flows if you selected every available port and multicast group, since each transmitter sends traffic to every selected multicast group (one to many). For more information about this pattern, see *"Multicast Backbone Traffic Pattern" on page 493*.

- The *Pair* pattern results in as many flows as there are pairs.

  The number of pairs is equal to the number of ports X number of multicast groups. For more information about this pattern, see *"Multicast Pair Traffic Pattern" on page 495*.

- In the *Fully-meshed* pattern, each multicast-enabled port involved in the test is both a transmitter as well as receiver.

  Unlike the other patterns for multicast, only the receiver ports that are part of the multicast groups are used for transmitters. This can potentially yield fewer flows. However, if there are more receivers than transmitters in the *Backbone* or *Pair* patterns, by comparison the *Fully-meshed* pattern produces more flows. For more information about this pattern, see *"Multicast Fully-meshed Traffic Pattern" on page 494*.

Both the *Backbone* and *Pair* pattern display all available ports for the transmitters of the multicast group and only the multicast-enabled ports for the receivers.

The more multicast-enabled ports and/or multicast groups that exist, the more potential there is to create a large number of flows and groups of flows.

Tip: To rapidly generate as many SmartFlow groups and flows involving as many multicast groups and ports as possible, use the *Backbone* pattern and select all of the transmitter ports and multicast groups.

For example, if the *Backbone* pattern is used with one transmitter and two multicast groups, two flows result, as shown in *Figure 19-15 on page 491*.

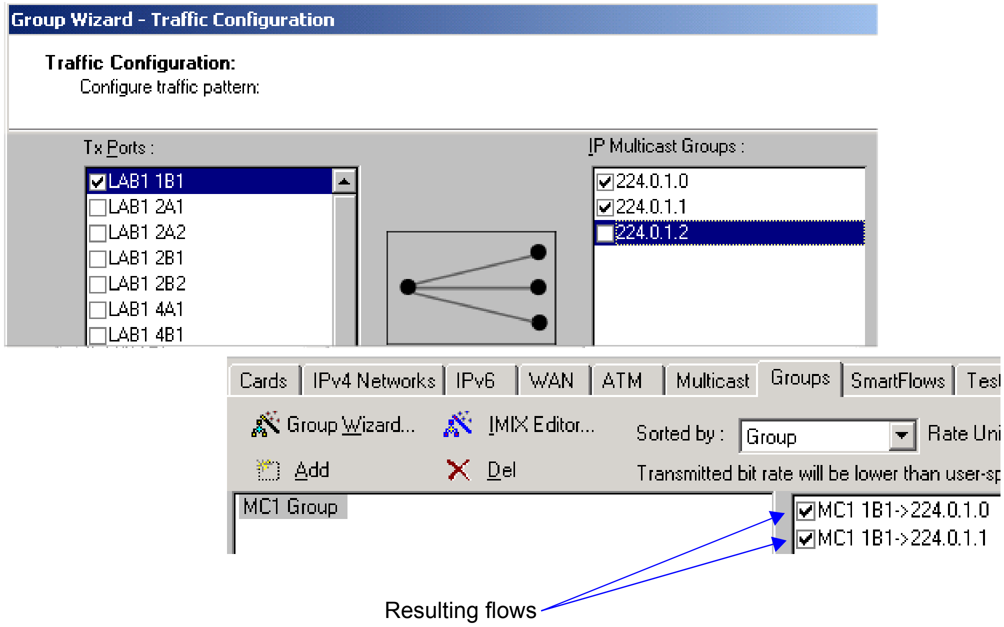

Figure 19-15. Example: Resulting Flows from One Tx and Two Groups, Backbone Pattern

If the *Pair* pattern is used with the same configuration of one transmitter and two multicast groups, the “same” two flows as present in the *Backbone* pattern also result, as illustrated in *Figure 19-16*.

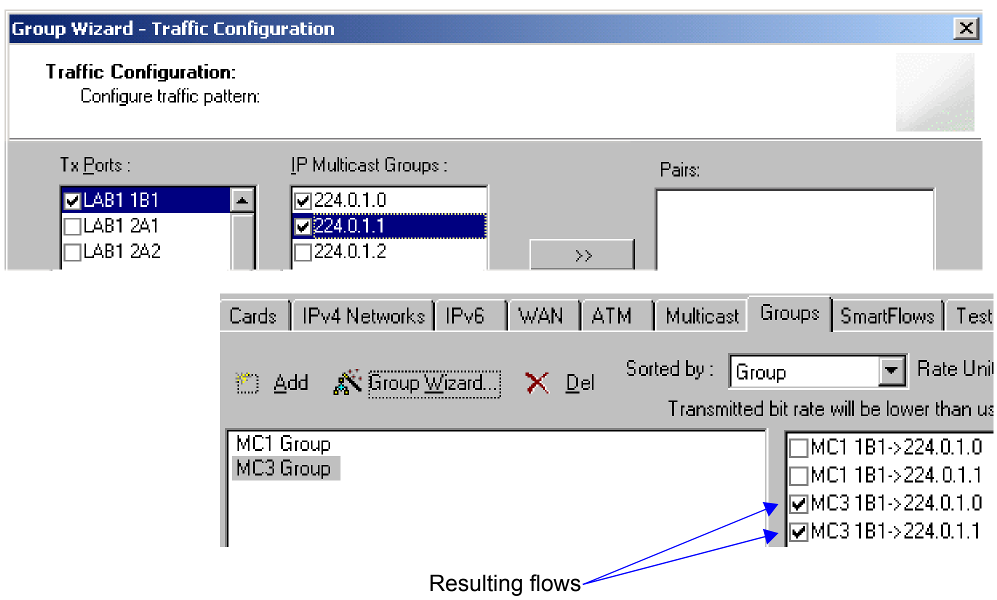

Figure 19-16. Example: Resulting Flows from One Tx and Two Groups, Pair Pattern

If the *Fully-meshed* pattern is used with the same two multicast groups, in this example more flows result than from the *Backbone* and *Pair* patterns, as shown in *Figure 19-17*. This is because the backbone and pair examples only used one transmitter, whereas the fully-meshed pattern by nature uses each receiver also as a transmitter in the selected multicast groups.

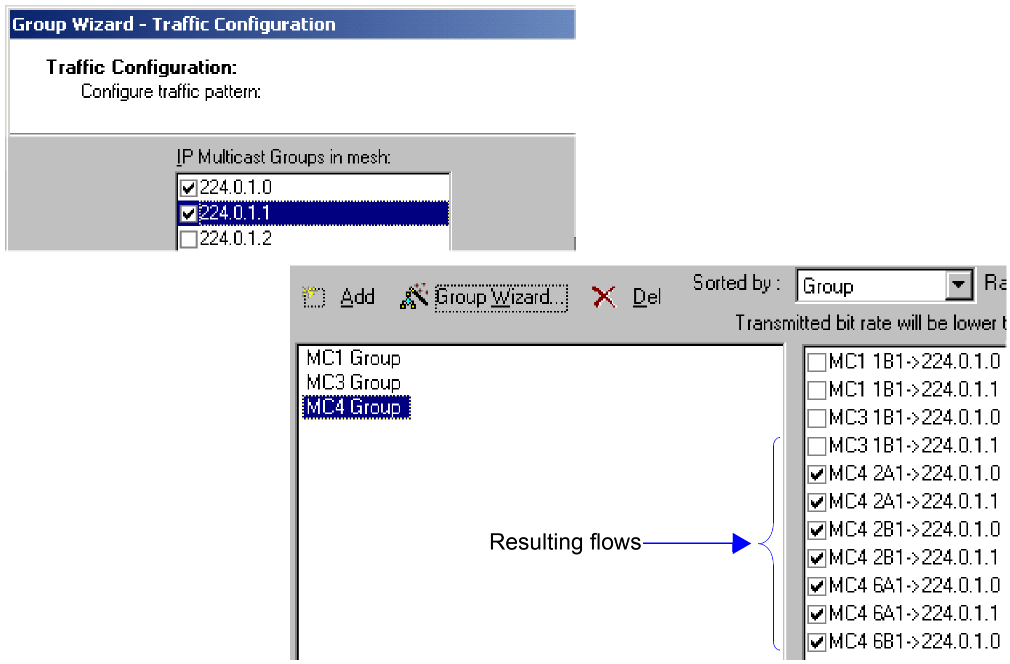

Figure 19-17. Example of Resulting Flows from Two Groups, Fully-Meshed Pattern

The *Fully-meshed* pattern with multicast traffic results in more flows than the *Backbone* or *Pair* patterns when there are more receivers than transmitters in the other patterns.

Multicast Backbone Traffic Pattern

The Group Wizard *Backbone* pattern for multicast traffic does not contain a traffic direction, since multicast groups are only a unidirectional pattern in which transmitters only send to receiver ports.

Let us say you used the configuration on the *Multicast* tab shown in *Figure 19-14 on page 489* (three multicast groups each with one untagged and tagged VLANs), with a traffic configuration consisting of one transmitter to the three multicast groups.

You set up multicast backbone traffic in the Group Wizard with one transmitter to multiple multicast groups, as shown in *Figure 19-18*.

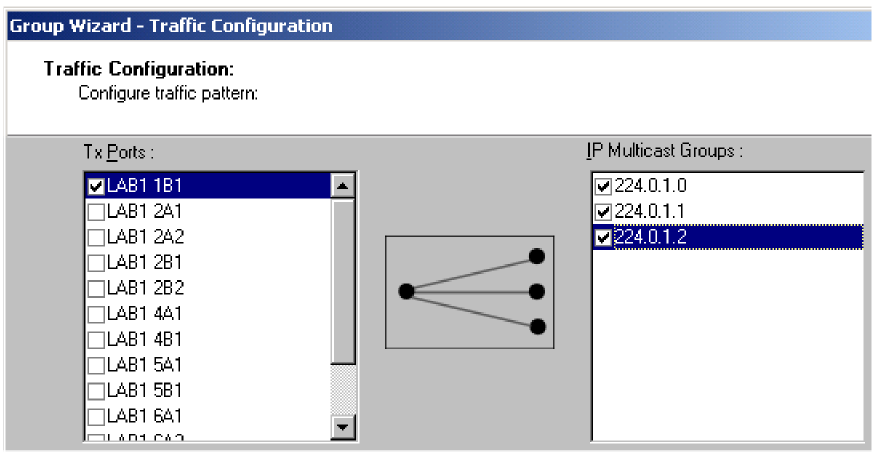

Figure 19-18. Group Wizard Backbone Pattern for Multicast Traffic

The transmitter ports do not have to be multicast-enabled, since they are not sending out IGMP join or leave requests.

This would result in the same port being the transmitter for multiple groups. *Figure 19-19* shows the SmartFlow group and SmartFlows that would result.

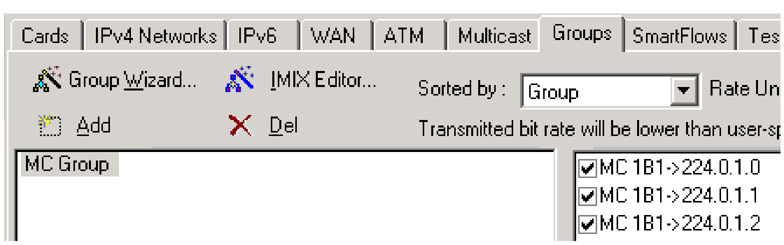

Figure 19-19. Resulting Flows from Backbone Pattern - One Tx to Multiple MC Groups

Multicast Fully-meshed Traffic Pattern

A fully-meshed pattern, in terms of multicast, means that every port within a multicast group is transmitting to and receiving from all ports in the group. The Group Wizard *Fully-meshed* traffic pattern is the best way to rapidly generate an even number of flows for each transmitting port. Use this pattern to simulate trunking, or if each port needs to receive the transactions of the other ports, such as in the London Stock Exchange.

The *Fully-meshed* traffic pattern for multicast traffic does not provide transmitter selection. Instead it displays only multicast groups and uses every receiver port that is displayed on the *Multicast* tab as a transmitter for each group selected. *Figure 19-20* shows the transmitting ports in the flows that correspond to the receiver ports on the *Multicast* tab.

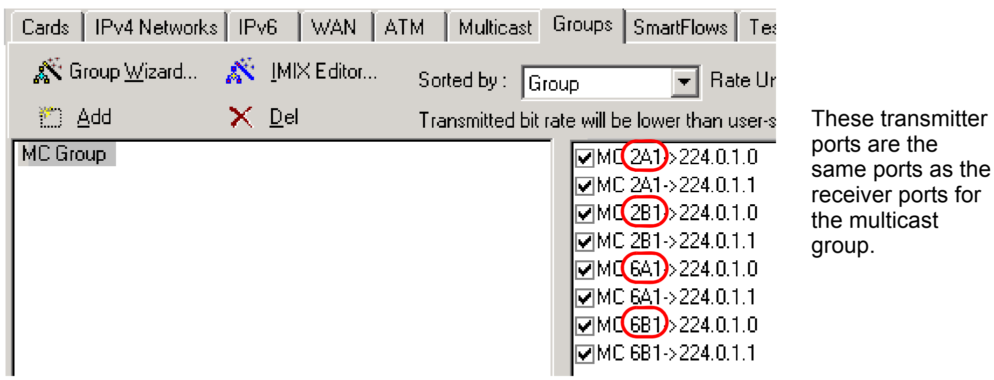

Figure 19-20. Resulting Flows from Mesh Pattern for Multicast

Because each transmitter is also a receiver, the number of flows created is double the number of receivers for each multicast group.

Multicast Pair Traffic Pattern

Each pair in the Group Wizard *Pair* traffic pattern for multicast traffic consists of one transmitter and one multicast group. You can make multiple selections from both the *Tx Ports* transmitter list and the *IP Multicast Groups* list, but the resulting pairs will each consist of one transmitter and one group. *Figure 19-21* shows an example to the *Pair* traffic pattern.

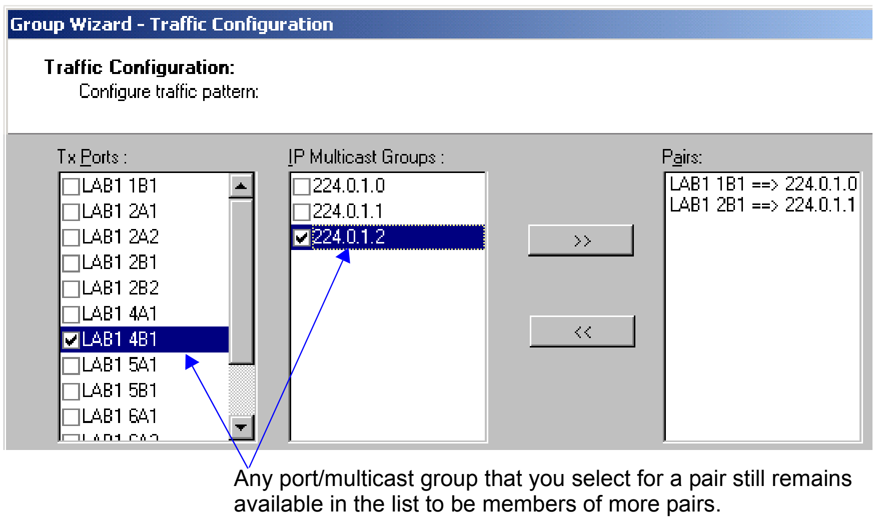

Figure 19-21. Group Wizard Pair Pattern for Multicast Traffic

The transmitter ports do not have to be multicast-enabled, since they are not sending out IGMP join or leave requests. Unlike unicast, if you select a pair of ports and multicast groups, they remain available in the list for any further pair selections. Once you select a transmitter port(s) and a multicast group(s), click the >> button to move the pair(s) into the *Pairs* column.

If you selected an uneven number of transmitter ports to multicast groups, SmartFlow reuses the ports or groups in a round-robin fashion to make pairs. *Figure 19-22 on page 496* illustrates this scenario and the pairs that result.

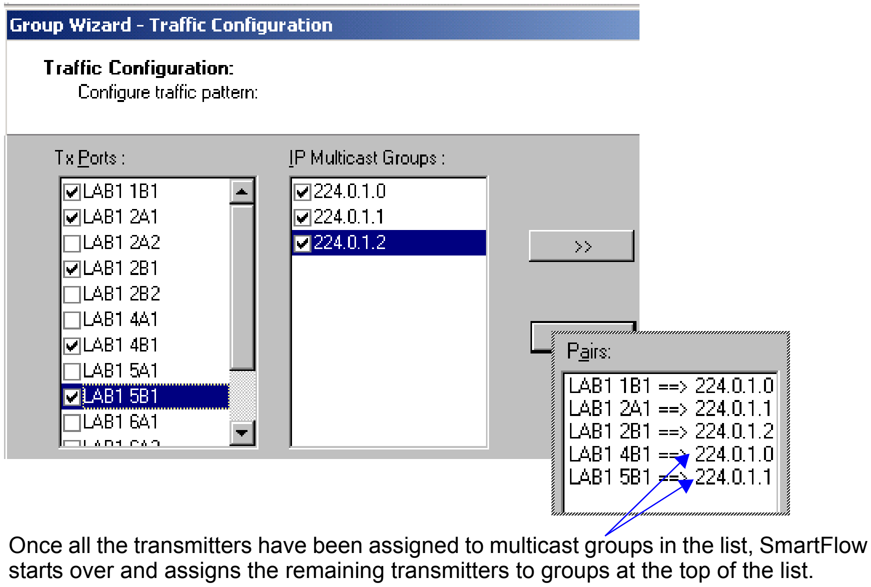

Figure 19-22. Uneven Transmitter/Multicast Groups in Pair Pattern

If you select more transmitters than multicast groups, once the transmitters are assigned to all of the multicast groups, SmartFlow starts over at the top of the list of multicast groups. It assigns the remaining transmitters to groups until each transmitter is assigned.

Interpreting IP Multicast Test Results

To track results by multicast group, it is recommended that when flows and groups are created, you set up a group for each multicast group that exists. This makes tracking results by multicast group easier to follow, since SmartFlow displays totals of each group together and then all flows together for each test iteration.

Figure 19-23 shows Summary results of a test that was set up so that each group of flows corresponds to a multicast group.

Jumbo (in uSec)

Name/Load	10.000%	20.000%	30.000%	40.000%	50.000%	60.000%	70.000%	80.000%	90.000%	100.000%
Total : Std. Dev.	27660.50	20394.40	16891.50	14743.10	17945.60	26268.60	26613.90	18052.90	20145.00	11358.50
MC1 Group	27940.10	20726.30	17816.40	16090.30	19374.50	26677.30	27492.10	19994.60	19333.30	11369.90
MC2 Group	25872.70	19099.80	15782.60	13584.70	16395.50	24820.30	25059.90	16249.50	15907.40	7021.50
MC3 Group	29582.20	21658.30	17047.00	14272.00	17811.40	27556.90	27486.60	16962.50	26444.40	15206.50

Figure 19-23. Jumbo Summary Results for a Mixed Class Test with Multicast Traffic

Multicast detail results are clearer if you set up one group for either each QoS level or each multicast transmitter. *Figure 19-24* shows how the detail results lists, by iteration, first group totals and then all flows in order of flow group.

Total for the test iteration.
Totals for each SmartFlow group, which in this case also represents a multicast group.
One SmartFlow.
An index of the port pair in that multicast SmartFlow.
Tx port name
Rx port name

Name	Packet Leng (bytes)	CRC Size (bytes)	Flow Loac (% util)	Tx Stream (streams)	Transmitte (packets)	Expected (packets)
Total	N/A	N/A	N/A	9	25320	101280
A DIFF0	N/A	N/A	N/A	3	8442	33768
A DIFF3	N/A	N/A	N/A	3	8439	33756
A DIFF6	N/A	N/A	N/A	3	8439	33756
A D0 1-1->224.0.1.0 (1: "SMB 1 1-1" --> "SMB 1 1-1")	124	4	1.111	1	2814	0
A D0 1-1->224.0.1.0 (2: "SMB 1 1-1" --> "SMB 1 1-2")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.0 (3: "SMB 1 1-1" --> "SMB 1 1-3")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.0 (4: "SMB 1 1-1" --> "SMB 1 1-4")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.0 (5: "SMB 1 1-1" --> "SMB 1 1-5")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.1 (1: "SMB 1 1-1" --> "SMB 1 1-1")	124	4	1.111	1	2814	0
A D0 1-1->224.0.1.1 (2: "SMB 1 1-1" --> "SMB 1 1-2")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.1 (3: "SMB 1 1-1" --> "SMB 1 1-3")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.1 (4: "SMB 1 1-1" --> "SMB 1 1-4")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.1 (5: "SMB 1 1-1" --> "SMB 1 1-5")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.2 (1: "SMB 1 1-1" --> "SMB 1 1-1")	124	4	1.111	1	2814	0
A D0 1-1->224.0.1.2 (2: "SMB 1 1-1" --> "SMB 1 1-2")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.2 (3: "SMB 1 1-1" --> "SMB 1 1-3")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.2 (4: "SMB 1 1-1" --> "SMB 1 1-4")	124	4	1.111	1	2814	2814
A D0 1-1->224.0.1.2 (5: "SMB 1 1-1" --> "SMB 1 1-5")	124	4	1.111	1	2814	2814
A D3 1-1->224.0.1.0 (1: "SMB 1 1-1" --> "SMB 1 1-1")	124	4	1.111	1	2813	0
A D3 1-1->224.0.1.0 (2: "SMB 1 1-1" --> "SMB 1 1-2")	124	4	1.111	1	2813	2813
A D3 1-1->224.0.1.0 (3: "SMB 1 1-1" --> "SMB 1 1-3")	124	4	1.111	1	2813	2813
A D3 1-1->224.0.1.0 (4: "SMB 1 1-1" --> "SMB 1 1-4")	124	4	1.111	1	2813	2813
A D3 1-1->224.0.1.0 (5: "SMB 1 1-1" --> "SMB 1 1-5")	124	4	1.111	1	2813	2813
A D3 1-1->224.0.1.1 (1: "SMB 1 1-1" --> "SMB 1 1-1")	124	4	1.111	1	2813	0
A D3 1-1->224.0.1.1 (2: "SMB 1 1-1" --> "SMB 1 1-2")	124	4	1.111	1	2813	2813
A D3 1-1->224.0.1.1 (3: "SMB 1 1-1" --> "SMB 1 1-3")	124	4	1.111	1	2813	2813

Figure 19-24. Jumbo Detail Results for Multicast Traffic

Note: Results are displayed per individual SmartFlow. However for multicast, the multiple port pairs that make up the multicast SmartFlow each appear on a separate line. The statistics for each of the port pairs of a multicast SmartFlow are repeated on each line. Thus, interpret multicast results line by line rather than try to add the transmission statistics.

IP Multicast Test Results by Group

To view results by multicast group instead of by SmartFlow, change the following line in the SmrtFlow.ini file:

```
Report=1
```

This causes SmartFlow to write results per multicast group to a file called **report.csv**, as shown in *Figure 19-25*.

15	Name	Packet Le	CRC Size	Flow Loac	Tx Strear	Transmitte	Expected	Received	Lost	In-Sequen
16		(bytes)	(bytes)	(% util)	(streams)	(packets)	(packets)	(packets)	(packets)	(packets)
63	A 1-1->224.0.2.13 (2: "SMB 6 1-1"	124	4	0.667	1	1689	1689	1678	11	1677
64	A 1-1->224.0.2.13 (3: "SMB 6 1-1"	124	4	0.667	1	1689	1689	1678	11	1677
65	Multicast Group 224.0.1.0	N/A	N/A	N/A	1	1689	5067	5058	9	5055
66	SMB 6 1-1 (Tx): A 1-1->224.0.1.0	124	4	0.667	1	1689				
67	SMB 6 1-2 (Rx): A 1-1->224.0.1.0						1689	1686	3	1685
68	SMB 6 1-3 (Rx): A 1-1->224.0.1.0						1689	1686	3	1685
69	SMB 6 2-1 (Rx): A 1-1->224.0.1.0						1689	1686	3	1685
70	Multicast Group 224.0.2.0	N/A	N/A	N/A	1	1689	5067	5058	9	5055
71	SMB 6 1-1 (Tx): A 1-1->224.0.2.0	124	4	0.667	1	1689				
72	SMB 6 1-2 (Rx): A 1-1->224.0.2.0						1689	1686	3	1685
73	SMB 6 1-3 (Rx): A 1-1->224.0.2.0						1689	1686	3	1685
74	SMB 6 2-1 (Rx): A 1-1->224.0.2.0						1689	1686	3	1685
75	Multicast Group 224.0.2.1	N/A	N/A	N/A	1	1689	5067	5055	12	5052
76	SMB 6 1-1 (Tx): A 1-1->224.0.2.1	124	4	0.667	1	1689				
77	SMB 6 1-2 (Rx): A 1-1->224.0.2.1						1689	1685	4	1684
78	SMB 6 1-3 (Rx): A 1-1->224.0.2.1						1689	1685	4	1684
79	SMB 6 2-1 (Rx): A 1-1->224.0.2.1						1689	1685	4	1684
80	Multicast Group 224.0.2.2	N/A	N/A	N/A	1	1689	5067	5055	12	5052
81	SMB 6 1-1 (Tx): A 1-1->224.0.2.2	124	4	0.667	1	1689				
82	SMB 6 1-2 (Rx): A 1-1->224.0.2.2						1689	1685	4	1684
83	SMB 6 1-3 (Rx): A 1-1->224.0.2.2						1689	1685	4	1684
84	SMB 6 2-1 (Rx): A 1-1->224.0.2.2						1689	1685	4	1684

Figure 19-25. Results per Multicast Group in .csv File

The **report.csv** file first lists each flow in the test and then each flow by each multicast group.

Appendix A

Frame Size Limits

This appendix provides the frame size ranges (by card technology) that SmartFlow allows to run a test. It also includes some recommended minimum frame sizes for certain cards.

In this appendix:

Allowable Frame Size Ranges by Card Type

You can enter frame size ranges that are illegal for a technology and still run the test, provided that the range falls within the limits on frame sizes (without CRC) imposed by SmartFlow. For example, you can run jumbo (9KB) frames with some Gigabit cards. For a list of the minimum and maximum frame sizes allowed per card/module, refer to the *SmartBits System Reference* guide.

To have the capability to transmit jumbo-sized frames, select the **Allow jumbo frame transmission** checkbox on the *Options* tab. SmartFlow allows you to transmit frame sizes that fall within the ranges shown in *Table A-1 on page 501* for Ethernet cards and *Table A-2 on page 506* for non-Ethernet cards. When viewing this table, keep in mind the following information:

- CRC size can vary by technology:
  Ethernet: 4 bytes.
  POS: 2 or 4 bytes, depending on your configuration.
  WAN: 0, 2 or 4 bytes, depending on your configuration.
- Frame size limits for frames with VLAN tags usually increases by four bytes.
- As with frames with VLAN tags, the size limits for frames with MPLS tags include the 4-bytes per label added to the frame. MPLS allows up to 4 labels per frame, so a frame with four labels results in 16 bytes added to the size.
- For WAN all protocols, using a frame size smaller than 80 bytes is not recommended since the offered test load deviates from intended test load.

Notes:
- If a LAN-310x module sends jumbo (9996 bytes) or oversized (16,000) frames, the maximum number of streams allowed is reduced to 128 and the maximum number of multicast groups is reduced to 300.
- LAN-33xxA, XLW-372xA, and XFP-373xA modules: If you send frames between 9997 and 16,384 bytes (without VLAN tag), they are reported only on the *Port Errors* results tab as an *Oversized* error.

Use *Table A-1 on page 501* and *Table A-2 on page 506* together with *Table A-3 on page 508* to determine the frame sizes allowed by card type.

Table A-1. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Ethernet Cards

	Frame Size (bytes) (without CRC)			
	Minimum		Maximum	
Protocol	**Non-VLAN**	**VLAN Tagged**	**Non-VLAN**	**VLAN Tagged**
IPv4				
LAN-320x	52	56	65,532	65,532
LAN-33xx	52 or 70 if DI[1] enabled	56 or 74 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XFP-373x	52 or 84 if DI enabled	56 or 88 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XLW-372x	52 or 84 if DI enabled	56 or 88 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
ML-5710 ML-7710	52	56	1600	1600
LAN-310x	60 or 65 if DI enabled	56 or 69 if DI enabled	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled
IPv6				
LAN-33xx	72 or 90 if DI enabled	76 or 94 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XFP-373x	72 or 100 if DI enabled	76 or 104 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results

Table A-1. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Ethernet Cards (continued)

	Frame Size (bytes) (without CRC)			
	Minimum		Maximum	
Protocol	**Non-VLAN**	**VLAN Tagged**	**Non-VLAN**	**VLAN Tagged**
XLW-372x	72 or 100 if DI enabled	76 or 104 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
LAN-310x	72 or 85 if DI enabled	76 or 89 if DI enabled	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled
UDP-IPv4				
LAN-320x	60	64	65,532	65,532
LAN-33xx	60 or 78 if DI enabled	64 or 82 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XFP-373x	60 or 100 if DI enabled	64 or 104 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XLW-372x	60 or 100 if DI enabled	64 or 104 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
ML-5710 ML-7710	60	64	1600	1600
LAN-310x	60 or 73 if DI enabled	64 or 77 if DI enabled	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled
UDP-IPv6				
LAN-33xx	80 or 98 if DI enabled	84 or 102 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results

Table A-1. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Ethernet Cards (continued)

	Frame Size (bytes) (without CRC)			
	Minimum		Maximum	
Protocol	**Non-VLAN**	**VLAN Tagged**	**Non-VLAN**	**VLAN Tagged**
XFP-373x	80 or 116 if DI enabled	84 or 120 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XLW-372x	80 or 116 if DI enabled	84 or 120 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
LAN-310x	80 or 93 if DI enabled	84 or 97 if DI enabled	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled
TCP-IPv4				
LAN-320x	72	76	65,532	65,532
LAN-33xx	72 or 90 if DI enabled	76 or 94 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not traced in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XFP-373x	72 or 100 if DI enabled	76 or 104 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XLW-372x	72 or 100 if DI enabled	76 or 104 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
ML-5710 ML-7710	72	76	1600	1600
LAN-310x	72 or 85 if DI enabled	76 or 89 if DI enabled	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled

Table A-1. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Ethernet Cards (continued)

	Frame Size (bytes) (without CRC)			
	Minimum		Maximum	
Protocol	**Non-VLAN**	**VLAN Tagged**	**Non-VLAN**	**VLAN Tagged**
TCP - IPv6				
LAN-33xx	92 or 110 if DI enabled	96 or 114 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XFP-373x	92 or 132 if DI enabled	96 or 136 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XLW-372x	92 or 132 if DI enabled	96 or 136 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
LAN-310x	92 or 105 if DI enabled	96 or 109 if DI enabled	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled
ICMP				
LAN-33xx	60 or 78 if DI enabled	64 or 82 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XFP-373x	60 or 92 if DI enabled	64 or 96 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
XLW-372x	60 or 92 if DI enabled	64 or 96 if DI enabled	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	1518 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results
ML-5710 ML-7710	60	64	1600	1600
LAN-310x	60	64	1514 or 16,000 if Jumbo enabled	1518 or 16,000 if Jumbo enabled

Table A-1. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Ethernet Cards (continued)

	Frame Size (bytes) (without CRC)			
	Minimum		Maximum	
Protocol	**Non-VLAN**	**VLAN Tagged**	**Non-VLAN**	**VLAN Tagged**
IP/MPLS				
(LAN-33xx only)	56-68	N/A	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	N/A
UDP/MPLS				
(LAN-33xx only)	64-76	N/A	1514 or 16,384 if Jumbo enabled 9997–16,384 not tracked in results	N/A

1 DI = Data Integrity

Table A-2. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Non-Ethernet Cards

Technology		Minimum Frame Size (in bytes) (without CRC)		Maximum Frame Size (in bytes) (without CRC)	
	Protocol	**Non-VLAN**	**VLAN Tagged**	**Non-VLAN**	**VLAN Tagged**
ATM	IPv4	38 if VC Mux Routed 46 if LLC Routed 54 if VC Mux Bridged 62 if LLC Bridged	N/A	32,767	N/A
	UDP-IPv4	46 if VC Mux Routed 54 if LLC Routed 62 if VC Mux Bridged 70 if LLC Bridged	N/A	32,767	N/A
	TCP-IPv4	58 if VC Mux Routed 66 if LLC Routed 74 if VC Mux Bridged 82 if LLC Bridged	N/A	32,767	N/A
	ICMP	46 if VC Mux Routed 54 if LLC Routed 70 if LLC Bridged 62 if VC Mux Bridged	N/A	32,767	N/A
POS	IPv4 Card Exceptions:	50	N/A	65,532 for 2-byte CRC, 65,530 for 4-byte CRC	N/A
	POS-3504/5 POS-3510/11	42 or 60 if DI enabled	N/A	16,380	N/A
	POS-3518/19	42 or 84 if DI enabled	N/A	16,380	N/A
	UDP-IPv4 Card Exceptions:	50	N/A	65,532 for 2-byte CRC, 65,530 for 4-byte CRC	N/A
	POS-3504/5 POS-3510/11	50 or 68 if DI enabled	N/A	16,380	N/A
	POS-3518/19	50 or 92 if DI enabled	N/A	16,380	N/A

Table A-2. SmartFlow Allowable Illegal Frame Size Ranges (without CRC) for Non-Ethernet Cards (continued)

Technology		Minimum Frame Size (in bytes) (without CRC)		Maximum Frame Size (in bytes) (without CRC)	
POS	TCP-IPv4 Card Exceptions:	62	N/A	65,532 for 2-byte CRC, 65,530 for 4-byte CRC	N/A
	POS-3504/5 POS-3510/11	62 or 80 if DI enabled	N/A	16,380	N/A
	POS-3518/19	62 or 104 if DI enabled	N/A	16,380	N/A
	IPv6 POS-3504/5 POS-3510/11	66 or 80 if DI enabled	N/A	16,380	N/A
	POS-3518/19	66 or 104 if DI enabled	N/A	16,380	N/A
	UDP-IPv6 POS-3504/5 POS-3510/11	74 or 88 if DI enabled	N/A	16,380	N/A
	POS-3518/19	74 or 112 if DI enabled	N/A	16,380	N/A
	TCP-IPv6 POS-3504/5 POS-3510/11	86 or 100 if DI enabled	N/A	16,380	N/A
	POS-3518/19	86 or 124 if DI enabled	N/A	16,380	N/A
	ICMP Card Exceptions:	50	N/A	16,380	N/A
	POS-3504/5 POS-3510/11	50 or 68 if DI enabled	N/A	16,380	N/A
	POS-3518/19	50 or 92 if DI enabled	N/A	16,380	N/A
WAN	IP/Frame Relay IP/PPP UDP/Frame Relay	48	N/A	8192	N/A
	UDP/PPP	64	N/A	8192	N/A

Frame Size Recommendations

These recommendations only apply to certain cards and modules.

For ML-7710 and LAN-3101A/B cards with cyclic flows:
The maximum achievable line rate may vary depending on frame size and the cyclic flow configuration. To ensure maximum line rate, use a frame length of at least 144 bytes (without CRC).

Certain cards require a specific minimum frame size in order to ensure that the UDP or TCP checksum is correct. This requirement only applies to the cards in *Table A-3*. Cards other than those listed in this table are not affected by this requirement.

Refer to the *Table A-1 on page 501* for Ethernet cards and *Table A-2 on page 506* for non-Ethernet cards to determine the minimum and maximum frame sizes by card type that are allowed (but not necessarily recommended) within SmartFlow.

Table A-3. Recommended Frame Sizes

Card	Minimum Frame Size (in bytes) (without CRC)
LAN-3201B LAN-3201C	Protocols, without VLAN: **UDP-IPv4:** 62 **TCP-IPv4:** 74 Protocols, with VLAN: **UDP-IPv4:** 66 **TCP-IPv4:** 78
POS-3500A POS-3500B POS-3502A	Protocols: **UDP:** 52 **TCP:** 64
WAN-3445	Protocols: **IP, TCP, UDP:** 80

Appendix B

ESD Requirements

Spirent Communications manufactures and sells products that require industry standard precautions to protect against damage from electrostatic discharge (ESD). This document explains the proper process for handling and storing electrostatic discharge sensitive (ESDS) devices, assemblies, and equipment.

The requirements presented in this document comply with the EIA Standard, *ANSI/ESD S20.20-1999: Development of an Electrostatic Discharge Control Program* and apply to anyone who handles equipment that is sensitive to electrostatic discharge. Such equipment includes, but it not limited to:

- All electronic assemblies manufactured by Spirent Communications
- Discrete and integrated circuit semiconductors
- Hybrid microcircuits
- Thin film passive devices
- Memory modules.

Caution: Failure to comply with the requirements explained in this document poses risks to the performance of ESDS devices as well as to your investment in the equipment.

General Equipment Handling

Whenever you handle a piece of ESDS equipment, you must be properly grounded to avoid harming the equipment. Also, when transporting the equipment, it must be packaged properly. Follow the requirements below to help ensure equipment protection.

- Wrist straps must be worn by any person handling the equipment to provide normal grounding.
- The use of foot straps is encouraged to supplement normal grounding. If foot straps are used exclusively, two straps (one on each foot) should be used. Note that foot straps are only applicable in environments that use ESD flooring and/or floor mats.
- Hold ESDS equipment by the edges only; do not touch the electronic components or gold connectors.
- When transporting equipment between ESD protected work areas, the equipment must be contained in ESD protective packaging. Equipment that is received in ESD protective packaging must be opened either by a person who is properly grounded or at an ESD protected workstation.

- Any racks or carts used for the temporary storage or transport of ESDS equipment must be grounded either by drag chains or through direct connection to earth ground. Loose parts that are not protected by ESD-safe packaging must not be transported on carts.

Workstation Preparation

The ideal setup for working with ESDS equipment is a workstation designed specifically for that purpose. *Figure B-1* illustrates an ESD protected workstation. Please follow the requirements listed below to prepare a proper ESD protected workstation.

- The ESD ground must be the equipment earth ground. Equipment earth ground is the electrical ground (green) wire at the receptacles.
- An ESD protected workstation consists of a table or workbench with a static dissipative surface or mat that is connected to earth ground. A resistor in the grounding wire is optional, providing that surface resistance to ground is $\geq 10^5$ to $\leq 10^9\ \Omega$.
- The workstation must provide for the connection of a wrist strap. The wrist strap must contain a current limiting resistor with a value from $\geq 250\text{K}\ \Omega$ to $\leq 10\text{M}\ \Omega$.
- ESD protective flooring or floor mats are required when floor-grounding devices (foot straps/footwear) are used or when it is necessary to move in between ESD protected workstations when handling ESDS equipment.

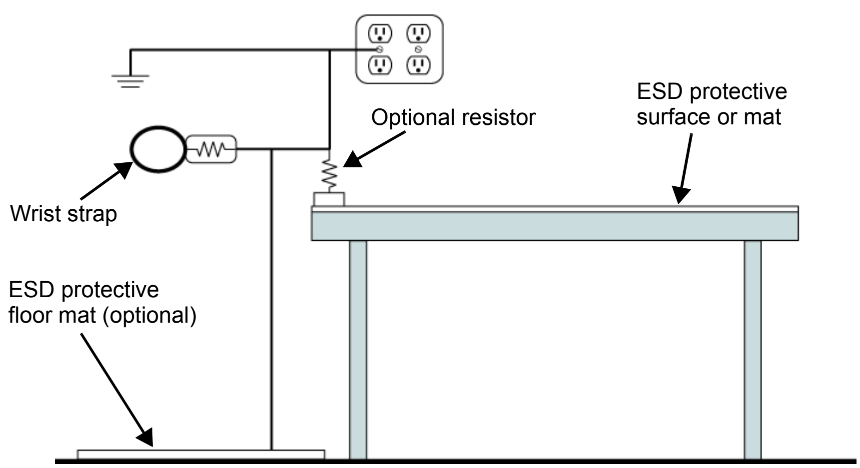

Figure B-1. ESD Protected Workstation

Note: The equipment needed for proper grounding is available in ESD service kits, such as the ESD Field Service Kit available from Spirent Communications (P/N 170-1800). Additional information on ESD can be found on the following website: http://www.esda.org/aboutesd.html

Appendix C

Fiber Optic Cleaning Guidelines

Spirent Communications manufactures and sells products that contain fiber optic components, including fiber optic transmitters and receivers. These components are extremely susceptible to contamination by particles of dirt or dust, which can obstruct the optic path and cause performance degradation. To ensure optimum product performance, it is important that all optics and connector ferrules be kept clean.

This document presents guidelines for maintaining clean fiber optic components. Spirent Communications recommends that these guidelines be followed very closely.

Caution:

- Failure to comply with the guidelines explained in this document poses risks to the performance of fiber optic-based devices, as well as to your investment in the equipment.
- Whenever you handle a piece of equipment that contains fiber optic components, you must be properly grounded to avoid harming the equipment. Refer to the Appendix in this document titled *ESD Requirements* for more details on ESD.

Cleaning Guidelines

To ensure the cleanliness of fiber optic components, follow the guidelines below:

- Use fiber patch cords (or connectors if you terminate your own fiber) only from a reputable supplier. Low-quality components can cause many hard-to-diagnose problems during an installation.
- Dust caps are typically installed on fiber optic components to ensure factory-clean optical devices. These protective caps should not be removed until the moment of connecting the fiber cable to the device. Ensure that the fiber is properly terminated, polished, and free of any dust or dirt. Also make sure that the location of installation is as free of dust and dirt as possible.
- Should it be necessary to disconnect the fiber device, reinstall the protective dust caps.
- If you suspect that the optics have been contaminated, alternate between blasting with clean, dry, compressed air and flushing with methanol to remove particles of dirt.

SPIRENT
Communications

Glossary

This glossary describes the terminology, abbreviations, and acronyms introduced in this guide, as well as terms based on telecommunications technology and standards.

ADPCM

Adaptive differential pulse mode modulation. A voice compression algorithm used by the G.726 codec to convert analog signals to digital information. It takes frequent samples of the sound, and predicts the next sample based on the amplitude and spectrum of the previous data. It uses an adaptive predictor (generally backward) to predict the spectrum of quantities of the next sample. It uses adaptive filters, based on the past accuracy of recent samples, to quantitize the difference between the prediction and the real signal. This difference is then used to update the predictor and quantizer for the next sample.

ADPCM is used to send sound on fiber-optic long-distance lines. One of the characteristics of note regarding this method is that it is relatively insensitive to bit errors (compared to pulse code modulation).

See also G.726.

Asymmetric Port Throughput Test

An additional level of testing that contains extra functionalities. Unlike the standard test, the asymmetric test is able to apply two sets of throughput parameters to two different groups of ports. Throughput is reported asymmetrically - and twice - for two sets of ports that are transmitting bi-directionally.

Autonegotiation

An optional standard function defined in the IEEE 802.3u Ethernet specification. Autonegotiation enables device interfaces to achieve the best possible mode of operation over a link segment, by providing a means to exchange information about capabilities (such as speed and duplex) over the link.

Autonegotiation with 10/100Mbps Ethernet links uses the MII (Media Independent Interface) registers of the PHY transceiver with the interface. There are 32 PHY registers altogether. However, only a few are involved in autonegotiation. Gigabit Ethernet links do not use PHY registers.

Autonomous system

A collection of networks under the control of one administrative authority (e.g., UUNET, Qwest, MichNet, etc.) that implement the same routing policy. The assignment of autonomous systems is strictly controlled. Abbreviation is AS.

The networks within an autonomous system communicate routing information to each other using Interior Gateway Protocol (Interior Gateway Protocol). Autonomous systems share routing information with other autonomous systems using Border Gateway Protocol (Border Gateway Protocol).

Bellcore

Bell Communications Research. One of the regional operating companies formed when AT&T split in 1984 into seven regional Bell operating companies (RBOCs). It developed technologies such as ISDN, ADSL, and SONET.

BGP

Border Gateway Protocol. A routing protocol used in TCP/IP networks. It is used exclusively to connect various networks and thus is classified as an inter-domain routing protocol.

Border Gateway Protocol is an exterior gateway routing protocol that enables groups of routers (called autonomous systems) to share routing information so that efficient, loop-free routes may be established. BGP is defined in RFCs 1771, 1772, 1773, 1774, and others. Current version is BGP4, defined in RFC 1771.

Binary Search Mode

A search mode used in a throughput test. The binary search mode uses a binary (scaled) search to determine a rate between the last failed rate and the last successful rate.

Class of Service

Used to designate a forwarding device as capable of prioritizing data traffic.

Combo Search Mode

A search mode used in a throughput test. This search mode starts with the step search mode, stepping up the rate until the test fails. Then it performs binary searches.

CRC

Cyclic Redundancy Code (CRC) used in cyclic redundancy checking. Cyclic redundancy checking is a method of checking for errors in transmitted data. The sending port applies a 16- or 32-bit polynomial to a block of data to be transmitted and appends the resulting CRC to the block. The receiving port applies the same polynomial to the received data and compares the result with the result appended by the transmitting port. If the results match, the data has been received successfully. If not, the transmitting port can be notified to resend the block of data.

A 16-bit cyclic redundancy code can detect all single and double-bit errors and ensures detection of 99.998% of all possible errors. This level of detection assurance is considered sufficient for data transmission blocks of 4 kilobytes or less. For larger transmissions, use a 32-bit CRC.

Customized Detailed Reporting

Customized detailed reporting is available for every SmartFlow test. The purpose of customized detailed reporting is to filter specific test results based on one or more criteria.

Cyclic Flow

A cyclic flow is a SmartFlow consisting of one stream definition with incrementing fields. It is called cyclic because, after SmartFlow has varied the field value(s) a specified number of times, it resets them to a base value and then begins to increment again.

When viewing results, the percentage of overall traffic that the flow uses is what characterizes it as a cyclic or non-cyclic.

Default LSP

A default LSP that uses default settings for TSpec, explicit route, and session attribute.

DHCP

The Dynamic Host Configuration Protocol (DHCP) streams enabler obtains IP addresses, gateway, lease time, and subnet (network) dynamically from the DHCP server and updates non-cyclic flows. The non-cyclic flows are automatically populated in the GUI with information from the DHCP server. Subsequent SmartFlow tests use the updated information in the test traffic.

DHCP coexists with statically configured clients. This means that some of the flows in the test (or on a port) could be DHCP-enabled while others can be assigned statically. (All SmartFlow tests can be run with DHCP-enabled flows.)

DiffServ

Differentiated Services. A method of prioritizing traffic, where high priority traffic is sent through or around clogged areas ahead of low priority traffic. Diffserv renames the IP Type of Service (TOS) field to the DS byte, which carries the information about IP packet service requirements.

DS3

A digital transmission rate of a digital communications circuit based on DS0, whose transmission rate is 64 Kbps. DS3 carries 28 DS1 signals or 44.736 Mbps. This digital signal hierarchy is part of the North American Digital Hierarchy (NADH).

DUT

Device under test

Egress LSR or Egress LER

A router at the egress edge of an MPLS network that receives MPLS packets forwarded through the core of the MPLS network and converts them back into IP packets.

Explicit Route

A route that is specified as a sequence of hops rather than being determined solely by conventional routing algorithms on a hop-by-hop basis.

The explicit route object is used to specify the route RSVP path messages take for setting up LSPs. Loose (rely on the routing table to find a destination) or strict (specify the next directly connected router) routes may be specified.

This parameter can be edited in SmartFlow to customize a default LSP.

Flow

One or more frames from one or more sources to one or more destinations that are tracked as a single entity. The sources and destinations can be defined at the IP, TCP, or UDP level.

Forwarding Equivalence Class (FEC)

A group of IP packets which are forwarded in the same manner. An example of an FEC is a set of packets whose source and destination addresses are identical. FECs may be defined at different levels of granularity to specify forwarding treatment.

Frame

In SmartVoIPQoS, the logical group of binary bits that results when a codec converts an analog sine wave into digital information. It is the minimum unit to carry voice data. You can have multiple voice frames in a single IP packet.

It is also used synonymously with the term *packet* when not referring specifically to a voice frame contained within an IP packet.

EBGP

External BGP. Refers to Border Gateway Protocol when it is run between different autonomous systems. SmartFlow supports only EBGP rules. (Compare with IBGP.)

G.711

An audio voice coding technique (codec) used in normal telephony by the PSTN. It is now being used with emerging voiceover IP standards such as H.323. In G.711, the encoded voice is already in the correct format for digital voice delivery through PBXs or in the PSTN. It uses a full 64 Kbps bit rate. It is part of the G-series recommendations in the ITU-T standard. With G.711, IP packets go out at a fixed rate and size.

G.723.1

An audio voice coding technique (codec) that compresses speech or audio signal components at a very low bit rate: 5.3 and 6.3 Kbps. It is part of the G-series recommendations in the ITU-T standard.

The higher bit rate yields a somewhat higher quality of sound and is based on ML-MLQ technology. The lower bit rate provides additional flexibility for system designers and is based on Code Excited Linear Prediction (CELP) or vector quantization of excitation.

With G.723.1, IP packet sizes vary.

G.726

An audio voice coding technique (codec) that uses Adaptive Differential Pulse Mode Modulation (ADPCM) to encode a G.711 bitstream into two-, three-, or four-bit words that results in bit rates of 16, 24, 32, or 40 Kbps. At 32 Kbps, the MOS score is 4.1, which is often regarded as the standard for toll quality.

This codec uses frames that are 125 usec long. *See also* ADPCM.

G.729

An audio voice coding technique (codec) that describes CELP compression where voice is coded into 8-Kbps streams. Two variations (G.729 and G.729 Annex A) of this standard exist. These variations differ mainly in computational complexity. Both variations provide speech quality similar to 32-Kbps ADPCM. It is part of the G-series recommendations in the ITU-T standard.

Gateway

A point (or device) in a network that serves as an entry to another network. Contrasted with host nodes such as the computers of network users and servers, gateways control the network traffic at an ISP or within a company network.

In SmartFlow, this is the IP address of the destination router's port.

Group

A number of flows combined together for purposes of analyzing test results.

H.323

An umbrella standard for how audio, video (such as in videoconferencing), and data communications take place over the Internet, LANs, and intranets. It is defined by the International Telecommunication Union (ITU). It includes the G.723.1 audio codec.

IBGP

Internal BGP. Refers to Border Gateway Protocol when it is run within an autonomous system. Compare with EBGP.

GPS

Global Positioning System. A system that determines your exact position on the Earth and can lock to a constant time from anywhere in the world. Based upon 24 satellites orbiting at 11,000 nautical miles above the Earth, GPS receivers may be used to synchronize the internal timing of the testing devices. Using GPS timing, testing between multiple Spirent devices can result in very accurate devices even though those devices may be thousands of miles apart.

Hop

The passage of a data packet from one router to the next.
RSVP Hop: Carries the IP address of the RSVP-capable node that sent this message.

ICMP

Internet Control Message Protocol. An extension to the Internet Protocol, it is the standard error and control message protocol for Internet systems. It is used to communicate between a gateway and a source host and detects and generates messages about congestion, network errors, and timeouts related to IP.

ICMP messages usually contain information about simple exchanges such as time-stamp or echo transactions (used by PING), or about routing difficulties with IP frames. It is defined in IETF RFC792.

Ingress LSR or Ingress LER

A router at the ingress edge of an MPLS network that receives IP packets, converts them to MPLS packets, and forwards them into the MPLS network.

IGMP

Internet Group Management Protocol. It is the standard protocol for IP multicasting on the Internet. It is used to set up host memberships in a multicast group on a single network. The host informs its local router that it wants to receive messages with the address for a particular multicast group. Hosts that conform to level 2 of the IP multicasting specification must use IGMP. IGMP is defined in RFC 1112.

IMIX

IMIX stands for Internet mix. It is a deterministic way of simulating real network traffic according to frame size usage. Some studies indicate that Internet traffic consists of fixed percentages of different frame sizes. IMIX traffic contains a mixture of frame sizes in a ratio to each other that approximates the overall makeup of frame sizes observed in real Internet traffic.

IP

Internet Protocol. Works in conjunction with TCP and is usually identified as TCP/IP.

IP Precedence

The TOS byte is divided into three sub-fields: IP precedence (3 bits), IP TOS (3 bits), and 2 bits reserved for future use. IP precedence is a method of assigning priorities to packets.

ITU-T

International Telecommunications Union - the Telecommunication Standardization Sector. (Formerly known as the CCITT.) Located in Geneva, Switzerland, it is the primary international body that promotes cooperative telecommunications equipment and systems standards.

Jitter

The amount of variation in latency between packets in a flow. Low jitter is important in voice transmissions. In SmartVoIPQoS, this value is found in the jitter column of the VoIP QoS test detailed results. It is similar to latency standard deviation in SmartFlow.

Label

A packet header used by an MPLS label switch router to forward packets. The header format depends upon network characteristics.

Label Distribution Protocol (LDP)

The LDP protocol uses TCP for reliable transport of packets through the MPLS network. Basic LDP does not support traffic engineering. The term *LDP* can also generically refer to any or all label distribution protocols.

Label Edge Router (LER)

The provider edge device that performs initial packet processing and classification and applies the MPLS label. This is a generic term for referring to all ingress and Egress LERs.

Label Switched Path (LSP)

A dynamic or static path assigned between endpoints in an MPLS network.

Label Switch Router (LSR)

A router that performs MPLS forwarding through the core of an MPLS network.

Latency

The time interval between the transmission and reception of a frame.

MPLS label stack entry

The four fields comprising the MPLS label that is inserted into an IP label stack at the ingress edge of an MPLS network. The values of the four fields are recognized and read by the MPLS core routers (LSRs) to enable the IP packet to be forwarded quickly and properly (based on QoS settings) through the core of the MPLS network.

MTU

Maximum Transmission Unit. The largest packet size in bytes that a network can transmit. Messages that are larger than the MTU are divided into smaller packets before being sent. The minimum MTU for IPv4 is 68 bytes; the minimum MTU for IPv6 is 1,280 bytes.

Multicast

A method of transmitting text, audio, and video on the Internet or an internal network to multiple receivers. Multicasting is a more efficient way of sending traffic because the server sends a single multicast packet to multiple multicast group recipients simultaneously. Contrast this one-to-many transmission with a unicast transmission, which requires a separate connection for each source-destination pair.

Multicast packets only travel along paths that lead to recipients that have requested to be part of the multicast group (transmission), which also cuts down on traffic. It is the technology used for transmitting webcasts such as live, on-demand audio and video broadcasting. Multicast IP addresses that can be assigned to multicast groups are Class D addresses that range from 224.0.0.0 to 239.255.255.255.

Multiplexing

Sending multiple messages or streams of information simultaneously on a single physical channel in the form of a single, complex signal, and then recovering the separate signals at the receiving end.

Multiprotocol Label Switching (MPLS)

MPLS is a means of making IP forwarding connection-oriented through an MPLS network. At the edge of an MPLS network, a label stack is inserted between Layer 2 and Layer 3 protocols of an IP packet—as a shim protocol. This label stack makes Layer 2 features available to Layer 3 by using a shorthand for IP packet forwarding that supports sophisticated packet classification with high-rate data forwarding.

Next Hop

The node to send data packets to in order to get the packets closer to the destination. It is the next “leg” or point of a trip taken by a data packet from one router or intermediate point to another in a network. In TCP/IP networks, the number of hops a packet takes in reaching its destination (or hop count) is stored in the packet header. Packets with very large hop counts are discarded.

NTP

Network Time Protocol. A protocol on top of TCP that assures accurate local time-keeping for atomic clocks and radio at various locations on the Internet. Devices use this protocol to obtain the most accurate time from the clocks and radio. It can synchronize distributed clocks within milliseconds over long time periods.

PBX

Private Branch Exchange. A type of circuit-switching system that handles voice calls as well as data links. It can be used in data networks that can carry voice calls.

A telephone system within an enterprise that is owned and operated by the enterprise to handle both internal and external calls. Users share a number of external phone lines. PBXs eliminate the need for a separate line to the telephone company’s central office for each user. PBXs usually include a switchboard for a human operator.

PCM

Pulse code modulation. A binary digital method of transmitting analog data.

It samples the analog signal amplitude at regular time intervals and rounds off the amplitude at each sampling to the nearest predetermined level. This is known as quantization.

Penultimate Hop Pop (PHP)

The label at the top of the stack is removed by the upstream neighbor of the upstream neighbor to the egress LER. The egress LER requested this action by advertising implicit null label using LDP. The egress LER does not have to look up and remove the label.

PIM

Protocol-Independent Multicast. In multicast routing, this protocol consists of two modes: PIM-dense and PIM-sparse. Dense mode involves flooding the network with traffic. Sparse mode involves establishing a rendezvous point to which all members send packets.

POS

Packet Over SONET.

PPP

Point-to-Point Protocol. A high-speed WAN transport protocol used to serially transmit IP data traffic over SONET frames with a minimum of overhead. It uses Link Control Protocol (LCP) packets to configure the data link and provides Layer 2 data-link service.

PSQM

Perceptual Speech Quality Measurement. A voice quality scoring system according to the ITU-T P.861. SmartVoIPQoS maps the latency variation (jitter) values and packet loss measurements of voice flows to a corresponding PSQM voice quality. (PSQM scores are not measured directly.)

PSQM scores range from 0 to 6.5, with 0 being the best and 6.5 being the worst.

PSTN

Public Switched Telephone Network. The collection of interconnected voice-oriented public telephone networks across the world, both commercial and government-owned. It is also known as Plain Old Telephone Service (POTS).

PVC

Permanent Virtual Circuit. In frame relay and in ATM, a logical point-to-point connection that is similar to a leased data line. It is dedicated and used over long periods of time.

QoS

Quality of Service. Like Class of Service, QoS indicates that a forwarding device is capable of prioritizing data traffic.

Quantization

In voice coding/decoding and Pulse Code Modulation (PCM), a process used in converting the analog signal to a digital information. It consists of rounding off the instantaneous amplitude of the analog signal at each sampling to the nearest of several specific, predetermined levels. The number of levels is always a power of 2 -- for example, 8, 16, 32, or 64.

Resource Reservation Protocol with Traffic Engineering (RSVP-TE)

An MPLS signaling protocol (an LDP) that uses the message acknowledgement of RSVP for reliable transport. The TE extension of RSVP-TE enables LSP establishment with traffic engineering. RSVP-TE is a soft-state protocol, requiring a refresh of its network topology information an average of every 30 seconds.

Route Flapping

Route flapping is any condition that causes a routing change that would result in a change in the BGP routing table (IP addresses of networks or groups of networks). Route flapping may be caused by a link outage, a router outage, or by command to withdraw routes. When the routing table changes, these changes cause the transmission of routing updates to be advertised to attached routers. Route flapping is one of the essential tests for testing BGP routers.

RTP

Real-time Transport Protocol. It is used to support real-time traffic such as voice and video and supports different types of payloads, including the ITU-T G.722 audio standard. It is the protocol used to carry the voice back and forth across a network.

The voice flows generated in SmartVoIPQoS are RTP flows.

SDH

Synchronous Digital Hierarchy. The international equivalent of SONET.

Session Attribute Object

Added to the PATH message by the ingress LER. This specifies LSP priority, preemption, and fast-reroute. It identifies the session. This value can be edited to customize an automated LSP.

SmartVoIPQoS

One of the applications in the SmartFlow Application Test Suite. The functions and interface are almost identical to SmartFlow in that you can set up multiple data flows of IP frames and test quality of service. However, SmartVoIPQoS allows you to define RTP flows that simulate voice traffic and contains a special voice over IP quality of service test. SmartVoIPQoS uses cards that simulate the IP encapsulated output of various codecs.

SONET

Synchronous Optical Network. A high-speed method of transport and international transmission protocol using optical fiber used for both voice and data. SONET is synchronous, meaning it requires an external clock source instead of inserting a clock bit in the transmission (as in asynchronous modes).

Standard deviation

Standard deviation is a statistical way of calculating how much the latency of a set of received packets varies from the mean. When individual latencies of received packets vary greatly from the mean, the latency has a high standard deviation. When latency does not vary greatly from the mean, the latency has a low standard deviation. It is similar to jitter.

Step Search Mode

A search mode used in a throughput test. This search mode increases the rate by the same percentage (step) added to the current rate.

SUT

System under test. Multiple devices under test configured as a network.

SVC

Switched Virtual Circuit. In frame relay and in ATM, a logical point-to-point connection that is set up for the duration of the connection and then disconnected. Similar to a telephone call, it is used for short periods of time.

TeraMetrics

The name of Spirent Communications' new open architecture for SmartBits modules and systems. A TeraMetrics module like the POS-3505As includes a Pentium-class processor running Linux Version 2.2. This provides on-board processing power sufficient for network-to-application layer performance testing at 1 terabit-per-second system speeds. It also makes possible support for third-party test applications.

Throughput

The maximum rate at which frames from flows and groups sent through a device can be sent without frame loss. There is a standard throughput test and an asymmetric throughput test.

TOS

Type of Service. An IP priority routing scheme. The 1-byte Type of Service field is in the header of an IP packet. Defined by RFC 1349, replaced by 2474 and updated by RFCs 3168 and 3260. Renamed by RFC 2474 to Differentiated Services, now called the Diff Serv Code Point (DSCP).

TTL

Time to Live. A threshold that is part of the IP packet representing the amount of time a multicast datagram has to traverse the network to its member hosts. When the packet arrives, the TTL is decremented. If TTL is zero or less, the packet is dropped.

Transit LSR

A labeled switch router that receives and forwards MPLS traffic.

TSpec Object

Contains link management configuration, including requested bandwidth plus minimum and maximum LSP packet size. This parameter can be edited in SmartFlow to customize an automated LSP.

VBR

Variable Bit Rate. One of the service types for transmitting traffic. Available in real-time or non-real-time versions.

VC

Virtual Circuit. A logical connection between two endpoints.

VLAN

A Virtual LAN. A configuration among network nodes based on a logical instead of physical connection.

WAN

Wide-area network. A data communications network that spans a broad geographical area and often uses transmission devices provided by common carriers. Examples of WANs are frame relay and ATM.

Index

D

N

SPIRENT
Communications