User's Manual

Ether.Genius
Ether.Sync
Ether.Giga

Gigabit Ethernet Testing Guide

ALBEDO Telecom

Copyright

Issue 6, 1/14

For any query or requirement regarding the *Ether.Genius*, *Ether.Sync* and *Ether.Giga Gigabit Ethernet Generator & Analyser* product line, contact with ALBEDO Telecom using the following contact details:

ALBEDO Telecom S.L.
C/ Joan d'Àustria 112
08018 Barcelona - Spain

E-mail: *support.telecom@albedo.biz*
Telephone: *+34 93 221 28 73*

ALBEDO Telecom - Joan d'Àustria, 112 - Barcelona - 08018 - www.albedotelecom.com

Table Of Contents

Chapter 1: Introduction 1

Important Notice 2
Warranty 2
Battery Safety 3
WEEE Notice 3

The Tester 3
Test Connectors 4
Platform Connectors 6

The Graphical User Interface 7

Running Tests 11

Upgrading the Unit 11

Chapter 2: Connection to the Network 13

Setting the Operation Mode 13

Connecting the Tester to the Network 16
Configuring the Connector 17
Using the SFP Ports 18
Choosing between One-way and Two-way Testing 20
Configuring Port A and Port B Auto-negotiation Parameters 21

Entering the Port Local Settings 24

Using the Traffic Reflector 26

Testing with an External Clock Reference 28

Using the Clock Reference Output 31

Chapter 3: Cable Tests 33

Electrical Test 33
Basic Principles of Ethernet Cable Wiring 33
Connecting the Tester 35
Running the Test 36

Measuring Optical Power 40

Chapter 4: Traffic Generation 41

Generation of Ethernet Traffic 41
Physical Layer Settings 42
Frame Settings 42
Configuring the Bandwith Profile 49
Choosing the Test Payload for Ethernet 52

Generation of IPv4 Traffic 55
Configuring the Physical and MAC Layers 56
Configuring MPLS 56
Configuring the Port Local Network Profile 60
Configuring the IPv4 Datagram 60
Setting the Bandwidth Profile 65
Choosing the Test Payload for IPv4 65

Event Insertion 65

Chapter 5: Basic Frame Analysis 69

Global Counts and Statistics 70
Frame Counts 70
Error Counts 72
Network Counts 73
Bandwidth Statistics 75
Frame Size Distribution 77
MPLS Statistics 77

Using the Network Search Capability 78

The LEDs Panel 79

The Event Logger 83
Configuring the Event Logger 84
Displaying Logs 85

BER Testing 87
Framed BER Tests 87
Physical Layer BER Tests 89

Chapter 6: Multi-Stream Analysis 95

Enabling and Disabling Filters 95

Configuring Filters 96
MAC Selection 97
C-VLAN and S-VLAN Selection 99
MPLS Selection 101
IPv4 Selection 103

IPv6 Selection 105
UDP Selection 107
Fixed Offset Selection 108

Per-Stream Counts and Statistics 110
Bandwidth Statistics 111
SLA Statistics 112

Chapter 7: Automatic Performance Tests 117

Performance Assessment with the RFC 2544 Test 117
Throughput 118
Latency 120
Frame Loss 121
Back-to-Back Frames 122
System Recovery 123
Configuration Pass / Fail Thresholds 124
Getting Test Results 125
Generation of RFC 2544 Result Reports 129

Performance Assessment with the eSAM Test 131
Bandwidth Profiles for Ethernet Services 133
Test Configuration 136
CIR Configuration Test 141
EIR Configuration Test 142
Policing Configuration Test 142
Performance Test 144
Generation of eSAM Result Reports 145

Chapter 8: Ping and Traceroute Tools 149

Ping 149
Internet Control Message Protocol 149
Test Configuration 150
Result verification 152

Traceroute 153
Test Configuration 154
Result verification 156

Chapter 9: IEEE 1588 Analysis 159

Ethernet Synchronization with IEEE 1588 159
Precedents: IP Synchronization with NTP 159
PTP Protocol Details 160

The Synchronization Mechanism 161
Protocol Encapsulation 162

IEEE 1588 Master and Slave Emulation 163
Configuration of the IEEE 1588 Master and Slave 164
Protocol State 167
Message Statistics 170
Delay Statistics 173
Slave Clock Status 177

Passive Monitoring 178

Chapter 10: Synchronous Ethernet Analysis 181

Introduction to Synchronous Ethernet 181
Ethernet Synchronization Messaging Channel 183
Synchronous Ethernet for the 1000BASE-T Interface 183

Synchronous Ethernet Frequency Measurement 186

Synchronous Ethernet Frequency Offset Generation 186

Operation in IP Through Mode 187

Chapter 11: Test Management 189

Generating Reports 189

File Management 191
Saving Configurations 192
Renaming Files 192
Deleting Files 192
Exporting Files to External Devices 193
Importing Configurations 193
Using the Embedded Web Server 194

Programming Tests 195

Using the System Menu 197
Using the Network 198
Installing Software Options 201
Using NTP for System Clock Synchronization 201

Using the Remote Control 203

Appendix A: Technical Specification 205

General 205

Operation Modes 205

Ethernet PHY **205**
Auto-Negotiation 205
Power over Ethernet 206
Synchronous Ethernet **206**
Operation 206
Analysis 206
Clock References **206**
Ethernet MAC **207**
MPLS **207**
IPv4 ..... **207**
Traffic Generator **208**
Bandwidth Profiles 208
Test Patterns and Payloads 208
Event Insertion **208**
Filter **208**
Ethernet Selection 208
MPLS Selection 208
IPv4 Selection 209
IPv6 Selection 209
UDP Selection 209
PHY Results **209**
Cable Tests 209
Auto-Negotiation 209
SFP 209
Power over Ethernet 209
Frame Analysis **210**
Ethernet Statistics 210
MPLS Statistics 210
IP Statistics 210
Bandwidth Statistics 210
SLA Statistics 210
BER 211
Network Exploration 211
Automatic Tests **211**
IETF RFC 2544 Test 211
eSAM Test 211
Port Loopback **211**
Ping and Trace-route **211**
PTP / IEEE 1588 ..... **212**
Results 212

Protocols **212**
User Interface **212**
Platform **213**

Appendix B: Common Issues and Solutions **215**

Chapter 1
Introduction

The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga are a family of handheld Ethernet traffic generators and analysers conceived to rate conformance and performance of native Ethernet networks or Ethernet services like Ethernet Private Lines (EPLs) or Ethernet Private LANs (EPLANs).

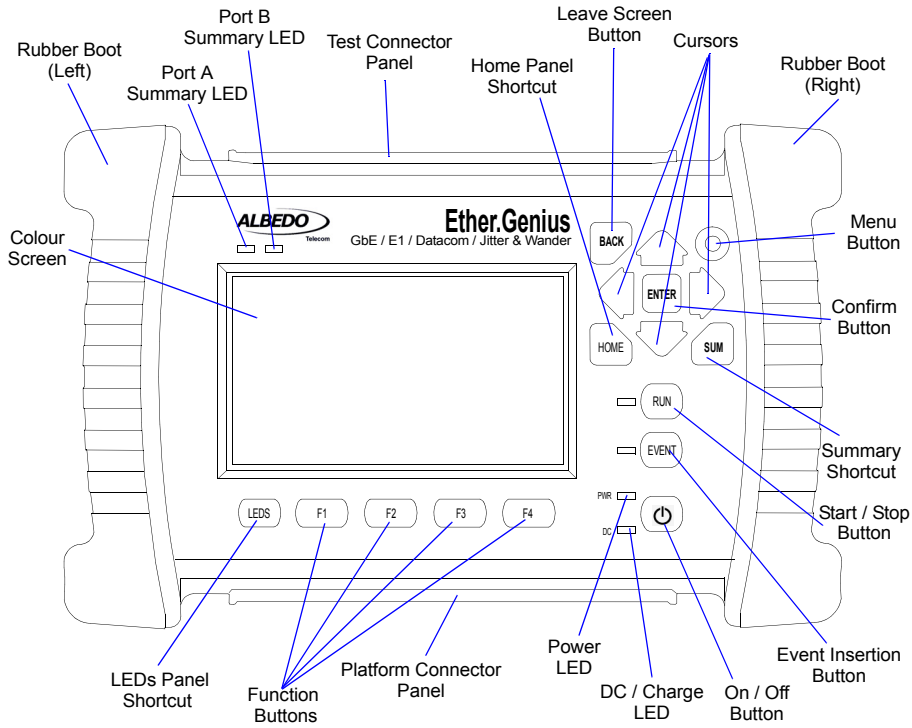

Figure 1.1: Ether.Genius front view. The tester presents results by means a colour screen and the LEDs. The configuration is performed through the keyboard.

- *Ether.Genius*: Is an hybrid TDM and Ethernet / IP generator and analyser. It comes with optional Ethernet synchronization testing capabilities. Ether.Genius is suitable for testing in environments where packet switching has not totally replaced legacy circuit switching technology, like in some cellular networks.
- *Ether.Sync*: This is an equipment with advanced Sync-E test and measurement capabilities like ESMC decoding or real-time MTIE / TDEV testing in Ethernet interfaces. Ether.Sync also comes with multistream traffic generation and analysis, RFC 2544 and eSAM (ITU-T Y.1564).
- *Ether.Giga*: Is a tester specifically designed for multistream Ethernet traffic generation and analysis over electrical and optical interfaces up to 1 Gb/s. Ether.Giga includes RFC 2544 and eSAM (ITU-T Y.1564) tests and advanced quality of service testing.

Ether.Genius / Ether.Sync / Ether.Giga have an external DC input but they also have internal batteries. This makes these testers suitable both for laboratory applications and field applications which require versatile and reliable operation.

Within your Ether.Genius / Ether.Sync / Ether.Giga test kit you will find the following items:

- One Ether.Genius / Ether.Sync / Ether.Giga test unit.
- One AC/DC adapter with a power cord specific for your country.
- One Carrying bag.
- Two Cat. 5e cables with RJ-45 connectors certified for operation at 1 Gb/s rates.
- Two SFPs for connection to optical interfaces (if ordered).
- Two MMF or SMF cables to be used with the SFPs (if ordered).
- Two coaxial cables with BNC male connectors (if ordered, Ether.Genius only).
- One datacom cable set (if ordered, Ether.Genius only).
- One CD-ROM with user documentation.
- One printed copy of this user manual (if ordered).

Check with your distributor the availability of other optional items for your Ether.Genius / Ether.Sync / Ether.Giga unit.

1.1. Important Notice

Operation, manipulation and disposal warnings for your Ether.Genius / Ether.Sync / Ether.Giga unit are listed below.

1.1.1. Warranty

The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga are supplied with a warranty that includes replacement of damaged or faulty components in the terms and period described in the ordering information. This Warranty does not apply to:

1. Product subjected to abnormal use or conditions, accident, mishandling, neglect, unauthorized alteration, misuse, improper installation or repair or improper storage.
2. Product whose mechanical serial number or electronic serial number has been removed, altered or defaced.
3. Damage from exposure to moisture, humidity, excessive temperatures or extreme environmental conditions.
4. Damage resulting from connection to, or use of any accessory or other product not approved or authorized by ALBEDO Telecom.
5. Product damaged from external causes such as fire, flooding, dirt, sand, weather conditions, battery leakage, blown fuse, theft or improper usage of any electrical source.

1.1.2. Battery Safety

The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga testers contain built-in batteries, improper use of which may result in explosion. Do not heat, open, puncture, mutilate, or dispose of the product in fire. Do not leave the device in direct sunlight for an extended period of time, which could cause melting or battery damage.

1.1.3. WEEE Notice

This product must not be disposed of or dumped with other waste. You are liable to dispose of all your electronic or electrical waste equipment by relocating over to the specified collection point for recycling of such hazardous waste. For more information about electronic and electrical waste equipment disposal, recovery, and collection points, please contact your local city centre, waste disposal service, or manufacturer of the equipment.

1.2. The Tester

Interaction with Ether.Genius / Ether.Sync / Ether.Giga is based on a high resolution colour screen, different kinds of status LEDs, and a keyboard. These are the keyboard elements:

- *Cursors*: Enable navigation through the graphical user interface. Including menus, keyboards and configuration lists. To leave a menu or configuration list, you can use the left arrow. In menus, the right arrow enters in the lower level menu or list.
- *ESC*: Leaves the current panel (menus, lists, and special panels).
- *ENTER*: In menus, enters in the lower level menu or list. In a configuration list or a keyboard panel, it selects the current item and leaves.
- *HOME*: Shortcut to the *Home* panel. From any menu, list, or special panel, it returns directly to *Home*.
- *SUMMARY*: Displays the *Summary* panel. If the Summary panel is already shown, it returns to the previous panel.

- *LEDS*: Displays the *LEDs* panel. If the LEDs panel is already shown, it returns to the previous panel.
- *MENU:* This key is reserved for future applications.
- *EVENT:* Starts (and sometimes stops) the event insertion. The exact way this button works depends on the actual event insertion mode. For example, if single event inversion is configured, each time you press *EVENT* a new event will be inserted but if the current insertion mode is configured to continuous insertion, you will need to press *EVENT* during event insertion to stop the action.
- *RUN:* This button starts / stops a new test. Some results (LEDs, some analogue values) are available without an ongoing test. Most of the configuration is blocked during a test execution. A test may need a short period of time to start after the *RUN* key has been pressed. It may also require some time to finish when *RUN* is pressed a second time. This condition is shown by a transient orange colour in the LED associated to the *RUN* key.
- *Function keys (F1, F2, F3, F4)*: These keys do not have a fixed purpose. Their associated action depends on the panel being displayed.
- *On/Off key*: If the tester is in off status, push to switch it on. If the tester is on, use this key to switch it off (long push).

There are four LEDs (PWR, DC, Port A summary, Port B summary). Their description is given below:

- *PWR*: Displays the current tester on / off status. The green colour is displayed under normal operation conditions. Orange and red are shown to indicate a low battery load.
- *DC*: This led is lit when the DC input is connected. Orange indicates a charging batteries status and green means that the internal batteries are ready.
- *Port A / Port B Summary*: These LEDs provide a permanent indication of the current input signal (or signals) status. The LEDs summarize the Port A and Port B information given by the event LEDs. If any event LEDs for a test port is in ‘red’ status, the port summary led will be set to ‘red’. If any event LED is ‘orange’ but there is no ‘red’ event, the summary led will be set to ‘orange. The ‘green’ colour is used when no events are found in the input signal. Finally, the LED is switched off when the port is disabled.

1.2.1. Test Connectors

Ether.Genius / Ether.Sync / Ether.Giga are connected to the DUT / SUT through the test connector panel. Ports and elements included in this panel are described in the following list:

- *RJ-45 Port A*. This is the primary 10/100/1000BASE-T port for Ethernet transmit and receive.
- *RJ-45 Port B*. This is the secondary 10/100/1000BASE-T port for Ethernet transmit and receive. This port is identical to the RJ-45 Port A in appearance but it pro-

vides only a subset of the features available for port A. Port B supports monitor and loopback operation but it does not include traffic generation.

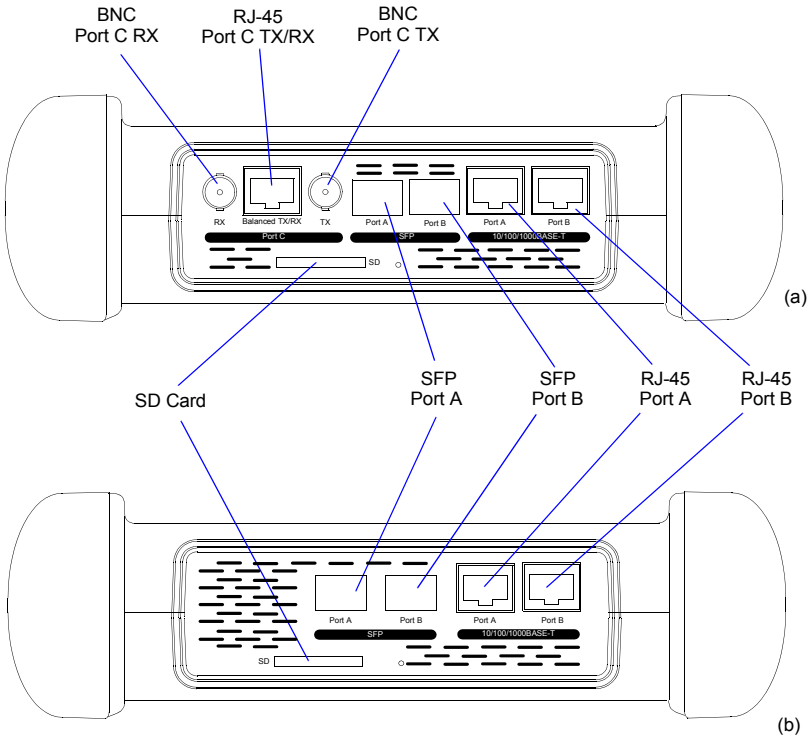

Figure 1.2: Test connector panel. Connection to the DUT / SUT is done in this panel: (a) Ether.Genius, Ether.Sync. (b) Ether.Giga

- *SFP Port A*. This port is used to connect the tester to the network through an optical interface with the help of an SFP module.
- *SFP Port B*. This port is used to connect the tester to the network through an optical interface with the help of an SFP module. This port is identical to the RJ-45 Port A in appearance but it provides only a subset of the features available for port A. Port B supports monitor and loopback operation but it does not include traffic generation.
- *SD Card*: Slot for SD Cards. These cards can be used as external storage devices.

- *BNC Port C RX* (Ether.Genius and Ether.Sync only): Unbalanced 75 Ω input. This input is used to analyse 2048 kb/s signals. It is used as a clock input port as well.
- *BNC Port C TX* (Ether.Genius and Ether.Sync only): Unbalanced 75 Ω output. This output is used to generate 2048 kb/s signals. It is used as a clock output as well.
- *RJ-45 Port C TX/RX* (Ether.Genius and Ether.Sync only): Balanced 120 Ω input / output. This interface is used to generate and analyse 2048 kb/s signals. This interface is used as a is used as a clock input / output as well.

1.2.2. Platform Connectors

There is a connector panel specifically devoted to the platform ports. This panel includes capabilities like remote control and external device connection. A more detailed description is given below:

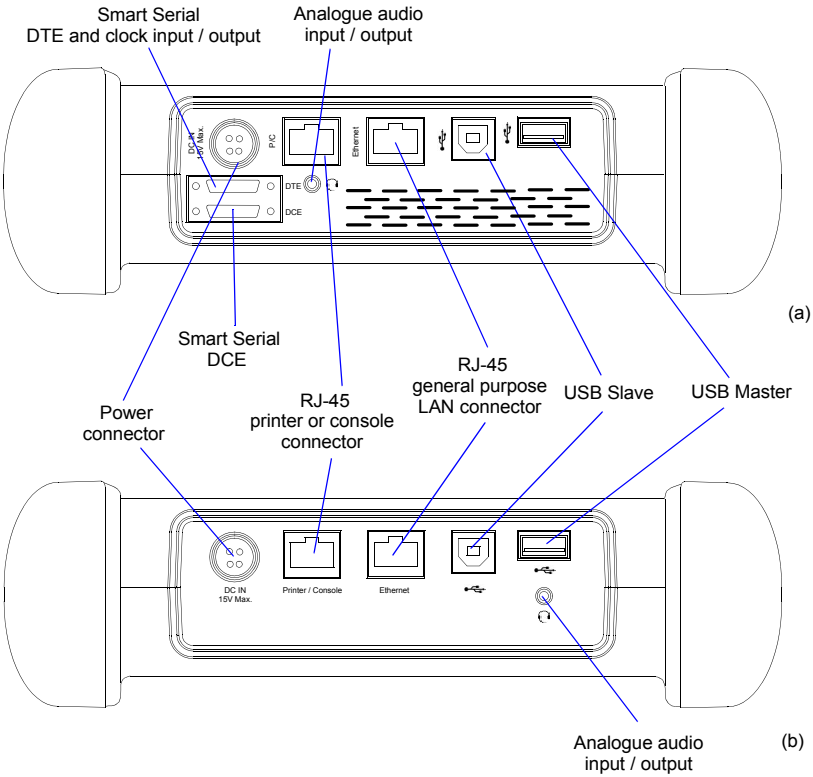

Figure 1.3: Platform connector panel. This panel includes connectivity to USB devices, Ethernet networks and power source: (a) Ether.Genius, Ether.Sync. (b) Ether.Giga.

- *Power connector*: The input must be 12 V DC, 4 A. A suitable external AC/DC adapter for your country is provided with the tester.
- *RJ-45 printer or console*. Console connector. This interface is prepared for connecting a serial printer.
- *USB Slave*. Use a USB cable with Slave type connector (Type B, *Device*) for this port. Currently this port enables connection of a PC to the tester and access to the internal tester file system.
- *USB Master*: Use a USB cable with a Master type connector (Type A, *Host*) for this port. Currently this port is used for software upgrades and connection of external storage devices.
- *RJ-45 general purpose LAN connector*: This is the platform Fast Ethernet connector (10/100BASE-T). It is used for remote management of the test unit or to access to the configuration and report files through a web interface.
- *Analogue audio input / output* (Ether.Genius and Ether.Sync): It is a 2.5 mm audio jack for connecting external speakers and microphone.
- *Smart Serial DTE* (Ether.Genius and Ether.Sync only): Universal datacom interface for the DTE. It supports V.11 / X.24, V.24, V.35, V.36, EIA-530, EIA-530A. This interface is used as a clock input for synchronization applications as well. The DTE connector can be considered as a test connector rather than a platform connector.
- *Smart Serial DCE* (Ether.Genius and Ether.Sync only): Universal datacom interface for the DCE. It supports V.11 / X.24, V.24, V.35, V.36, EIA-530, EIA-530A. The DCE connector can be considered as a test connector rather than a platform connector.

1.3. The Graphical User Interface

The Ether.Genius / Ether.Sync / Ether.Giga graphical user interface is based in a 480 x 272 colour screen and a set of keys attached to the front panel. Some of these keys have a permanent purpose but the specific function for some other keys depend on the context.

The keyboard and the screen allow the user setting configuration values, starting tests and displaying results. The user is always aware of the current status of the received signal through the Softleds shown on the left. The Softleds are always displayed and they work even when there is no test running. On the top side of the screen there is a

header zone which contains information about the current tester status (date, time, tests running, event insertion active) and an identifier for the currently displayed panel.

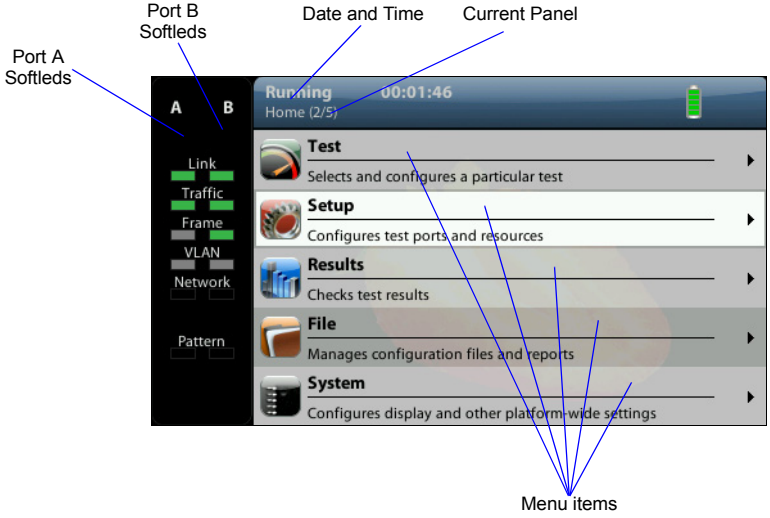

Figure 1.4: Ether.Genius / Ether.Sync / Ether.Giga Home panel.

Most of the graphical user interface panels are menus containing a variable number of items. All the menus are available from the *Home* panel. Users can press the HOME button at any time to go to the *Home* panel. The *Home* panel contains the following menu items:

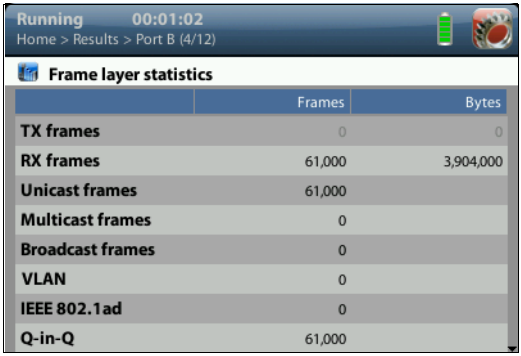

Figure 1.5: Ether.Genius / Ether.Sync / Ether.Giga test results represente as a counter list.

- *Test*: Contains configuration items related with general test configuration like delayed test settings, test setup, performance objectives, event insertion, and report configuration.

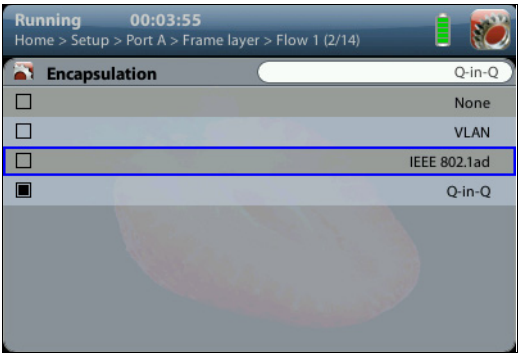

(a)

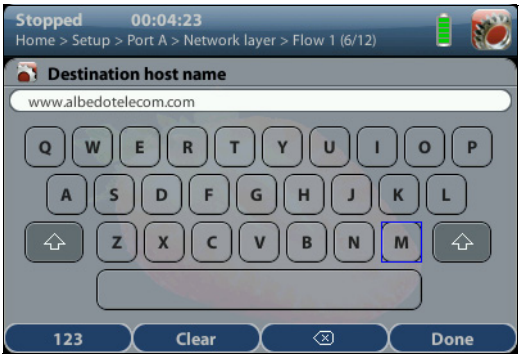

(b)

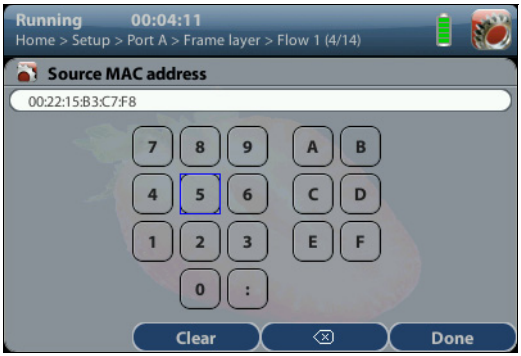

(c)

Figure 1.6: Different kinds of configuration panels: (a) Selection list, (b) Alphanumeric keyboard, (c) Numeric keyboard.

- *Setup*: Provides access to test resource configuration. For Ether.Genius / Ether.Sync / Ether.Giga, the setup submenus contain configuration of Port A, Port B, Port C and Datacom port.
- *Results*: This item enables the user to browse test results. Most of them are not available if a measurement has not been previously started but there are some exceptions to this rule like the LEDs.
- *File*: File management menus. Includes configuration and report file management. Files can be deleted, copied, exported or imported.
- *System*: Provides platform management tools. For example language selection, screensaver configuration and others.

There are two special panels as well. These are the *Summary* panel and the *LEDs* panel. The *Summary* provides some details about the current configuration and results. The *LEDs* panel gives extended the information about the received signal status already given by the Softleds. Both the *Summary* panel and the *LEDs* panel can be displayed at any moment by pressing the *SUM* and *LEDS* buttons.

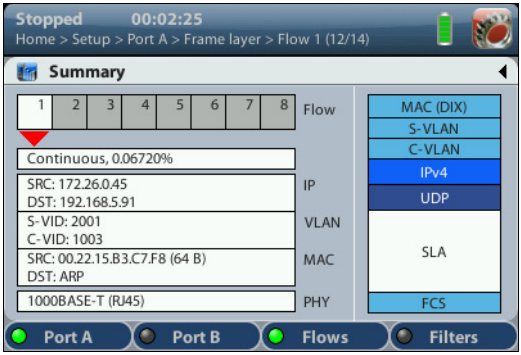

(a)

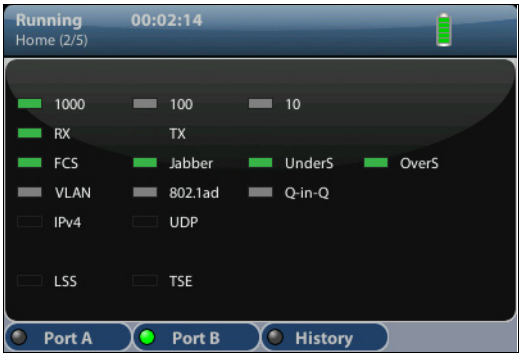

(b)

Figure 1.7: Special panels: (a) Summary panel, (b) LEDs panel.

Menus and submenus are organized in a tree. The root of the tree is the *Home* panel and the leaves are configuration or result panels. Results are usually presented in a list or a table. If all results cannot be simultaneously displayed, then the user is allowed to use the cursors up and down to browse the list.

Configuration panels are usually selection lists. Sometimes you can select only one simultaneous item in the list and sometimes selection of several items at the same time is possible. Keyboards are available if selection through lists is not possible. There is one keyboard for numeric settings and one for alphanumeric settings.

1.4. Running Tests

Most of the results provided by Ether.Genius / Ether.Sync / Ether.Giga are not available until you start a test. This section provides a high level description of the procedure to follow to configure your unit, start a test and review the results.

1. Configure the tester to send / receive signals in the right operation mode and through the right ports. Connect it to the network.
2. Program the test start time and duration with the help of the *Program* menu (within *Test*) or start the test immediately by pressing RUN.
   Note: Most of the configuration is blocked when there is an ongoing test.
3. Wait for the test to finish or press RUN to finish immediately.
4. Check the test results in the *Results* menu.
   Note: All test results are upgraded in real time as the test progresses. That means that is not really necessary to wait for the test to finish to check current results.

1.5. Upgrading the Unit

The test unit software can be upgraded with the help of a USB memory stick. Before proceeding with the upgrade copy the ALBEDO software to the root directory in the memory stick. The file name of the upgrade package must not be modified. The USB must have a FAT32 file system.

Once the USB memory stick is ready. Follow this procedure to install the new software:

1. Switch the unit off
2. Press HOME and ENTER simultaneously and, without releasing the keys, press the On / Off button.
3. Now, keeping all three the keys pressed, wait until you hear a beep. Then release the keys.
   The ALBEDO Software Installer is loaded and executed. An informative panel displays the Ether.Genius / Ether.Sync / Ether.Giga software version number found in the storage device.
4. Press ENTER to continue with the installation process.

5. Select *Install* or *Upgrade*. *Install* regenerates all the software in the unit even if it is up to date. Upgrade regenerates only the software that has changed since the last upgrade. Use *Install* (F2 key) if you need to recover the unit after operation failure due to corrupted software. Use *Upgrade* (F1 key) otherwise.
6. Confirm your previous selection by pressing ENTER or cancel with ESC.
7. Wait for the installation process to finish.
   Note: The full process may take a few minutes.
   Note: Do no disconnect the unit or remove the USB memory stick during installation.
8. Press ENTER to close the Software Installer and finish the installation process. The unit will be automatically restarted. The new software will be loaded.

Chapter 2
Connection to the Network

Ether.Genius / Ether.Sync / Ether.Giga are equipped with two identical 1 Gb/s RJ-45 ports and two 1 Gb/s SFP ports. The RJ-45 ports are used for connection to Ethernet electrical interfaces. The SFP ports are normally used for optical connections. Each RJ-45 / SFP interface constitutes a single logical port. These ports are labelled as Port A and Port B. The Ether.Genius model includes an extra port (Port C) for TDM signal generation and analysis.

Port A and Port B do not share the same generation and analysis capabilities. Port A is a full featured Ethernet port but Port B is conceived as a secondary testing port only. For this reason, Port B does not include advanced traffic generation. Port B can loop frames / packets toward their origin if configured to do so. It also responds to pings (ICMP echo request message) and other basic protocols like ARP. Analysis capabilities of Port A and Port B are similar, the only exception being the cable test.

This chapter describes how to connect the tester to the network and how to configure it to receive and send signals. The general, high level procedure to do that is:

1. Configure the Port A / Port B, including global and port specific operation modes and generation / analysis properties.
2. Connect the test cables to the network. Use the electrical or optical ports depending on the particular network properties.
3. Traffic generation does not start until you start a test with RUN. Most of the results are not available neither.

2.1. Setting the Operation Mode

As Gigabit Ethernet interfaces, Port A and Port B are independent but they share the same global operation mode. That means that, for example, if you configure Port A to be an IP endpoint port, then Port B will become an IP endpoint port as well. In any other sense, Port A and Port B are allowed to have a different configuration.

To configure the Ether.Genius / Ether.Sync / Ether.Giga global operation mode follow these steps:

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Select *Mode* to enter in the mode selection menu
3. Choose between: *Ethernet endpoint*, *IP endpoint*, *IP Through*, *Ethernet cable test* or *L1 endpoint*. Confirm by pressing ENTER.
   Note: Some other operation modes may be available depending on the hardware / software options available in your test unit.

Table 2.1: Global Ethernet Operation Modes

Mode	Description
Ethernet endpoint	If the tester is configured in *Ethernet endpoint* mode, then it emulates an Ethernet network terminating point. In this mode the generator sends a test signal made up of Ethernet frames to the DUT / SUT and the analyser receives new Ethernet frames from the DUT / SUT. The received signal may be the same test signal once it has been transmitted through the DUT / SUT. This mode is used for continuity / transparency testing, BER testing and performance (SLA, eSAM, RFC 2544) measurements in Ethernet interfaces. You can use this operation mode if you have IP over Ethernet but you don't care about the network protocol used in the network.
IP endpoint	If the tester is configured in *IP endpoint* mode, then it emulates an IP network terminating point. In this mode the generator sends a test signal made up of IPv4 frames to the DUT / SUT and the analyser receives new IPv4 / IPv6 frames from the DUT / SUT. The Received signal may be the same test signal once it has been transmitted through the DUT / SUT. This mode is used for continuity / transparency testing, BER testing and performance (SLA, eSAM, RFC 2544) measurements in IP / Ethernet interfaces.
IP through	The *IP Through* mode is suited for bidirectional intrusive monitoring. The signal from the Port A receiver is forwarded to the Port B transmitter. An equivalent operation is performed on the signal received on Port B. As it passes through the test unit, the test signal is analysed and statistics about it are collected and recorded.

Table 2.1: Global Ethernet Operation Modes

Mode	Description
Ethernet cable test	This is the correct mode to check the UTP / FTP / STP cable transmission parameters like the wire map, skew and MDI / MDIX port status. It there is any cable fault like an open or a short circuit, it is detected and an estimated distance to the fault is displayed. If this mode is enabled, port B is forced to a *Link* status and port A is set to an special *Cable test* status that is specifically used for cable tests.
L1 endpoint	This is a mode specifically conceived for physical layer BER tests. When the equipment is configured in this mode is unable to generate user-configurable frames but it can still generate and analyse the PCS codes required for BER testing at L1.

There is a port specific operation mode that is complementary to the global operation mode. Both the global and the port specific operation modes are combined to determine which tests and which capabilities are available at any moment.

Table 2.2: Port Modes

Mode	Description
TX / RX	Both transmission and reception are enabled in the port. The transmitter is connected to the internal test traffic generator. This is the port mode to be used most of the times for testing in endpoint mode. The *TX / RX* mode is available for Ethernet Port A only.
Monitor	The port is configured in promiscuous monitoring mode and the port transmitter is disabled. All kinds of traffic generation are disabled in *Monitor* mode but the equipment can still reply to some ARP and ICMP requests for technical reasons. Use this mode if you want to get statistics about the network traffic, including traffic from remote Ether.Genius, Ether.Sync or Ether.Giga units but you don't want to disturb the network with any test traffic internally generated by the unit.
Loopback	The port receiver is connected to the transmitter so that part or all the received frames are sent towards the origin. This port mode is used guarantee the continuity of the test payload or pattern to in two-way tests.

Table 2.2: Port Modes

Mode	Description
Cable test	This is a mode specifically used for cable tests. It is used if the global operation mode is set to *Ethernet cable test*. The *Cable test* port mode disables transmission of Ethernet frames or IP packets. It also stops any layer 2 or layer 3 analysis. The *Cable test* port mode is available for port A only.
Link	This is a mode specifically used for cable tests. It is used if the global operation mode is set to *Ethernet cable test*. The *Link* test mode disables transmission of Ethernet frames or IP packets. It also stops any layer 2 or layer 3 analysis. The only function of the *Link* port mode is to supply link to enable *Port A* to measure the all cable parameters. The *Link* port mode is available for port B only.
Disabled	Both the port transmitter and receiver are disabled. Use this mode if you are not going to use the corresponding port and you want to extend the operation time with batteries to the maximum.

To configure the Ether.Genius / Ether.Sync / Ether.Giga port specific operation mode follow these steps:

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific configuration menu.
3. Select *Mode* to enter in the mode selection Menu
4. Choose between: *TX / RX*, *Monitor*, *Loopback, Cable test, Link* or *Disabled*. Confirm by pressing ENTER.
   Note: Some other operation modes may be available depending on the global operation mode and the actual port you are configuring.

2.2. Connecting the Tester to the Network

The way you connect your tester to the network depends on the global operation mode. For example, if you set the tester to operate in *IP through* mode, you will be required to

connect the equipment in “transparent” mode, with the traffic going through the unit from Port A to Port B and Port B to Port A.

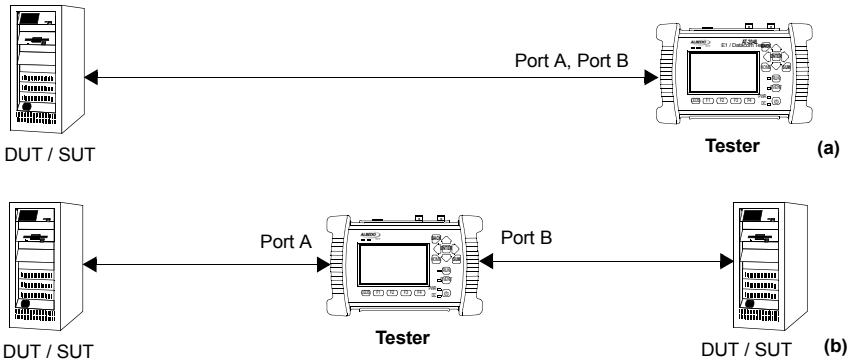

Figure 2.1: Basic connection setup for Ether.Genius, Ether.Sync and Ether.Giga testers: (a) Ethernet and IP endpoint operation modes. (b) Ethernet / IP through mode.

The operation mode also depends largely of the type of DUT / SUT. If you are connecting the equipment to an IP router, probably it does not make sense to configure the equipment in *Ethernet endpoint* mode because the router will be unable to recognise and forward traffic without an IP payload. Note that the opposite is not true. You may want to send IP traffic trough an Ethernet network and for this reason, the *IP endpoint* mode is compatible with Ethernet network testing. Many users (carriers and service providers) are offering Ethernet services like Ethernet Private Lines (EPLs) or Ethernet Private LANs (EPLANs) and they prefer to avoid making any decision concerning IP addressing. They also prefer to avoid generation of any IP stack protocol (ICMP, ARP, DHCP, PPP) within their administrative domains. For these users, the *Ethernet Endpoint* mode is the most appropriate.

2.2.1. Configuring the Connector

Both Port A and Port B have one Electrical (RJ-45) and SFP connector each. Which one is enabled at any moment is a user decision. Normally, the SFP is used to connect the equipment to an optical interface but nothing stops the user to attach a compatible electrical SFP module to Port A or Port B and carry out an electrical test with an SFP.

The procedure for configuring the connector (RJ-45 or SFP) is the following

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific configuration.
3. Go to *Physical* to enter in the physical settings configuration panel.
4. Select *Connector* to display the available options for the port connector.

5. Choose the right connector and confirm by pressing ENTER.
   Note that Port A and Port B do not need to be both optical or electrical at the same time. Ether.Genius / Ether.Sync / Ether.Giga are compatible with operation requiring conversion between optical and electrical transmission if all other operation conditions are met.

2.2.2. Using the SFP Ports

The SFP ports are the only choice available for optical tests. They can also be used for electrical tests if compatible SFPs are connected but this is usually not necessary due to the attached RJ-45 ports which require no adapters.

Table 2.3: Ethernet SFP Results

Result	Description
SFP present	Shows information about presence of an SFP in the current port..
Transceiver	Displays the current Ethernet interface. Supported interfaces are listed below: • 10BASE-T: Used for transmission at 10 Mb/s over two pairs of Cat. 3 UTP cable with range of 100 m. • 100BASE-TX: Used for transmission at 100 Mb/s over two pairs of Cat. 5 UTP cable with range of 100 m. • 100BASE-FX: Used for transmission at 100 Mb/s over two MMF in the 1310 optical window. This interface requires an special SGMII compliant SFP supplied by ALBEDO Telecom. • 1000BASE-T: Used for transmission at 1000 Mb/s over four pairs of Cat. 5e UTP cable with range of 100 m. • 1000BASE-SX: Used for transmission at 1000 Mb/s over two MMF operating in the 850 nm optical window. Ranges are usually a few hundred metres. This interface is supported by means an external SFP only. • 1000BASE-LX: Used for transmission at 1000 Mb/s over two MMF or SMF in the 1310 nm optical window. Ranges use to be a few kilometres. This interface is supported by means an external SFP only.
Vendor	If there is an SFP connected to the port, this field shows information about the vendor. This information is recorded within a memory in the SFP when it is manufactured.

Table 2.3: Ethernet SFP Results

Result	Description
Model	If there is an SFP connected to the port, this field shows information about the vendor. This information is recorded within a memory in the SFP when it is manufactured.
TX optical power	In compatible optical SFPs, this field displays the currently transmitted optical power expressed in dBm. Resolution for this result is 0.1 dBm but real result accuracy depends on the specific SFP module.
RX optical power	In compatible optical SFPs, this field displays the currently received optical power expressed in dBm. Resolution for this result is 0.1 dBm but real result accuracy depends on the specific SFP module.
Wavelength	In optical SFPs, this field displays the nominal wavelength used by the light source.

To display the SFP interface information follow these step sequence:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *SFP information* and check the *SFP present*, *Transceiver*, *Vendor*, *Model*, *TX optical power*, *RX optical power* and *Wavelength*.

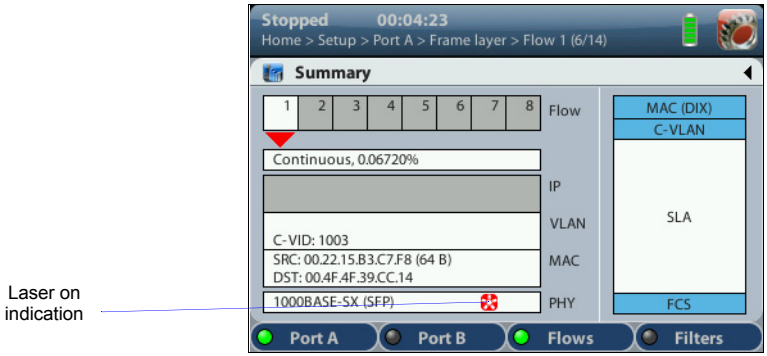

Figure 2.2: *Laser on* indication in the Ether.Genius / Ether.Sync / Ether.Giga Summary panel.

For security reasons, the optical transmitter is not automatically enabled when the equipment boots up. To switch the optical transmitter on, follow this procedure:

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific configuration.
3. Go to *Physical* to enter in the physical settings configuration panel.
4. Set *Laser* to *On* to enable the optical transmitter in the current port.
   Note: It is recommended to switch the optical transmitter once the testing has finished.

Current on / off status of the optical transmitter is always available from the *Summary* panel.

2.2.3. Choosing between One-way and Two-way Testing

Two-way tests measure network performance using a closed path in the network. There is usually a traffic reflector like Ether.Loop that returns the traffic towards the origin. The return path is not always the same that the forward path but the test traffic always finishes in the same port where it was originated. ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga support two-way tests in their Port A. Port B can be used as a local traffic reflector but installing an external reflector in a remote network location is also possible.

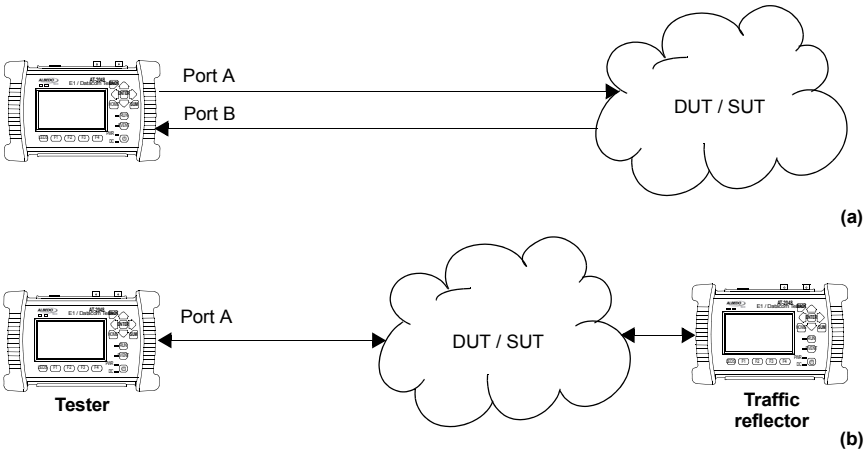

Figure 2.3: Connection setup for one-way and two-way tests: (a) One-way test, the test pattern leaves the tester in port A and it is received in port B. (b) Two-way test, the test pattern leaves in port A and it is received in the same port A.

One-way tests are done in open paths in the network. Traffic is generated in a test interface, transmitted through the network and analysed by a second test interface.

One-way tests provide a more accurate description of network performance than two-way tests but they require more difficult setups and have more requirements than two-way tests. Specifically, delay measurements require the generator and the analyser to use a common timing source. Delay variation measurements does not require the transmitter and the receiver to be synchronized but short term variations on the frequency offset between the transmitter and the analyser may also degrade the result accuracy.

Currently, Ether.Genius / Ether.Sync / Ether.Giga support one way measurements between Port A and Port B in the same unit. In this way, it is guaranteed that the generator and the analyser are using the same timing source. To switch between the one-way and two-way operation follow these steps:

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Select *Test method* to enter in the test method selection menu.
3. Choose between: *One-way (Port A-B)* or *Two-way (Port A-A)*. Confirm by pressing ENTER.
   Note: The *Test method* setting affects to all measurements that use SLA results as inputs. These include the RFC 2544 and eSAM automatic tests (See chapter 7).

2.2.4. Configuring Port A and Port B Auto-negotiation Parameters

Many Ethernet ports use auto-negotiation to negotiate speed, duplex operation and other parameters with the peer interface. Ether.Genius / Ether.Sync / Ether.Giga may or may not use auto-negotiation. If auto-negotiation is enabled, the user can decide whether to restrict the available bit-rates. If the user decides to disable auto-negotiation, then the bit rate is forced to a user configurable value.

Table 2.4: Ethernet Auto-negotiation Setup

Setting	Description
Enable	Enables or disables the standard Ethernet auto-negotiation procedure during the connection setup. Auto-negotiation sets the link bit rate, duplex mode, flow control mode without user intervention. Disable Auto-negotiation only if you know that the remote end does not support this procedure or if you want to check link operation without auto-negotiating.
1000-FD	Allows / disallows the interface to negotiate the 1000 Mb/s transmission rate. This is the only choice available if an optical SFP transceiver for 1000BASE-X is attached to the port.
100-FD	Allows / disallows the interface to negotiate the fast Ethernet speed (100 Mb/s) if you are using the RJ-45 port or a compatible electrical SFP for data transmission.

Table 2.4: Ethernet Auto-negotiation Setup

Setting	Description
10-FD	Allows / disallows the interface to negotiate the 10 Mb/s transmission rate if you are using the RJ-45 port or a compatible electrical SFP for data transmission.
Clock role	Sets the master / slave clock role in 1000BASE-T interfaces. The available settings for this field are listed below: • *Auto*: The master / slave role designation is automatic and in principle random. • *Master*: The synchronization master role is forced for the interface. This is the right configuration for frequency offset generation in Ether.Genius / Ether.Sync. • *Slave*: The synchronization slave role is forced for the interface. This is the right configuration value for frequency measurements over the interface. The clock role setting only makes sense if the equipment is allowed to negotiate the 1000BASE-T interface over a native RJ-45 port.

Under normal circumstances, the preferred link speed is the highest available. Use *1000-FD*, *100-FD* and *10-FD* settings if you want to analyse the link operation under sub-optimal circumstances. For e example use it to operate at 10 or 100 Mb/s in a link supporting 1000 Mb/s. The general procedure to configure the auto-negotiation in your Ether.Genius / Ether.Sync / Ether.Giga is as follows:

Table 2.5: Ethernet Auto-negotiation Results

Result	Description
Local	Displays the bit rate and duplex mode supported by the equipment Port A or Port B. If the current connector is the RJ-45, the supported bit rates are 10 Mb/s, 100 Mb/s and 1000 Mb/s (1000FD, 100FD and 10FD). If the connector is set to SFP the supported bit rate is 1000 Mb/s (1000FD).
Remote	Displays the bit rate and duplex mode supported by the remote device connected to Port A or Port B. It is one or several of 1000FD, 1000HD, 100FD, 100HD, 10FD, 10HD.
Current	Bit rate and duplex operation agreed during the auto-negotiation process. It is one (and only one) of the 1000FD, 100FD and 10FD set. If there is more than one compatible interface, the one with higher bit rate is preferred.

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific configuration.
3. Go to *Physical layer* to enter in the physical settings configuration panel.
4. Go to *Autonegotiation*.
5. Set *Enable* to *Yes* to configure the link speed using auto-negotiation or to *No* to disable auto-negotiation and force the link speed to a fixed value.
6. If you have enabled auto-negotiation in the previous step configure the allowed bit rates through the *1000-FD*, *100-FD* and *10-FD* menus. If auto-negotiation is disabled, set the *Forced bit-rate* to *10* or *100* from the *Physical layer* menu.
   Note: The 1000 Mb/s rate cannot be forced and it is only available through auto-negotiation due to IEEE 802.3 restrictions.

Once the tester has been connected to the network and the right connector type as been configured, follow these steps to check auto-negotiation results:

Figure 2.4: Albedo Ether.Genius / Ether.Sync / Ether.Giga auto-negotiation results panel.

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select *Port A or Port B* to enter in the port specific results.
3. Press *Auto-negotiation*
   The auto-negotiation status table for the current port is displayed.
4. Verify which bit rate / duplex modes combinations are supported by the local interface and the remote peer in the *Local* and *Remote* rows. Check the current speed and duplex mode in the *Current* row.
   Note: In IP / Ethernet through mode, auto-negotiation of each port depends on the auto-negotiation results from the second test port. For this reason, the current speed and duplex modes may be not the port optimum ones.

2.3. Entering the Port Local Settings

Ethernet devices do not normally require any configuration to work other than their factory-set MAC address but when the equipment is configured in *IP Endpoint* mode it behaves like any other IP host in the network. That means that the test unit should be able to receive traffic directed to it or transmit test and signalling traffic to other devices. To do that, users should assign an IP profile to each test port. Port A is full featured and it needs and IP address, network mask, gateway address and DNS address to operate. Port B does not need to generate traffic, it only receives traffic from other interfaces or responds to some requests. For this reason, Port B does not need a full IP profile.

Table 2.6: Local Ethernet / IP Profile

Result	Description
Use DHCP	Defines the procedure used to set the IP profile to the current port. The port IP profile is made up of an IPv4 address, a network mask, an optional gateway address and an optional DNS address. If *Use DSCP* is enabled, then the *Dynamic Host Configuration Protocol* (DHCP) defined in standard RFC 1531 will be used to set the IP profile in the current port.
Static IPv4 address	Is the 32 bit IPv4 address assigned to the local port in decimal, four-dotted format. Addressing scheme used here follows standards RFC 790 and RFC 791. The static addressing is used only if no dynamic address configuration like DHCP is used for the port.
Static IPv4 network mask	Subnet mask used to identify the network address bits and host address bits in the *Static IPv4 address*. The use of the subnet mask is defined in RFC 1219. The subnet mask is entered in decimal, four-dotted format and it must belong to the same network (same network bits) than the current port. The static addressing is used only if no dynamic address configuration like DHCP is used for the port.
Static IPv4 gateway	IP address corresponding to the network device used to send IP packets to external networks. The gateway address is configured in decimal, four-dotted format. The gateway address is used only by the port transmitter (not by the receiver). For this reason, this setting is available only in port A.

Table 2.6: Local Ethernet / IP Profile

Result	Description
Static IPv4 DNS server	IP address corresponding to the host used for domain name resolution. A DNS server allows the user to identify destinations by alphanumeric domain names rather than numeric IP addresses. The DNS address has to be entered in decimal, four-dotted format. The DNS address is useful only by transmitting IP packets. For this reason, this setting is available only in port A.
Leased IPv4 address	Current DHCP-assigned IP address in a decimal four dotted format. This is a read-only field that cannot be directly configured by users This setting makes sense only if *Use DHCP* is enabled.
Leased IPv4 network mask	Current DHCP-assigned network mask in a decimal four dotted format. This is a read-only field that cannot be directly configured by users This setting makes sense only if *Use DHCP* is enabled.
Leased IPv4 DNS server	Current DHCP-assigned DNS server in a decimal four dotted format. This is a read-only field that cannot be directly configured by users This setting makes sense only if *Use DHCP* is enabled. The DNS server is used only by the port transmitter (not by the receiver). For this reason, this setting is available only in port A.
Leased IPv4 gateway	Current DHCP-assigned default gateway in a decimal four dotted format. This is a read-only field that cannot be directly configured by users This setting makes sense only if *Use DHCP* is enabled. The gateway address is used only by the port transmitter (not by the receiver). For this reason, this setting is available only in port A.
Ethernet address	This is a read-only field that displays the factory MAC address assigned to the port. The traffic generator supports emulation of MAC traffic from sources other than the factory MAC, but the address displayed here is the real, unique MAC address used in default configuration.

In order to configure the IP profile in your Ether.Genius / Ether.Sync / Ether.Giga unit follow these steps:

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific configuration.
3. Enter in *Local profile*.
4. Decide whether you want to configure the local IP profile with the help of the DHCP protocol or you want to statically set these profiles by means the *Use DHCP* control.
   If you have enabled DHCP wait for the tester to get an IP profile from a DHCP server . If DHCP is not enabled, enter a valid *Static IPv4 address*, *Static IPv4 network mask*, *Static IPv4 gateway* and *Static IPv4 DNS address*.
   Note: For Port B, you only need to enter a valid IPv4 address and network mask .

2.4. Using the Traffic Reflector

You can configure Ether.Genius / Ether.Sync / Ether.Giga test ports in loopback mode so that the traffic they receive is forwarded toward their originator. This is very useful in many test and measurement applications where the traffic generator is waiting to receive the test traffic so that it can correlate the frames with the original ones.

Table 2.7: Local Ethernet / IP Profile

Result	Description
Loop mode	The loop mode determines which fields within the Ethernet frame or the IP datagram are swapped before forwarding. Is one of the following ones: • *Physical loop*: Loops frames without other alteration than pulse shape regeneration. This mode may cause problems if it is used in a bridged network. • *MAC loop*: Swaps source and destination MAC addresses before forwarding the frame. This is the correct mode to be used in bridged networks. • *IP loop*: It operates in the same way that the *MAC loop* mode but is swaps source and destination IP addresses as well. This is the correct mode to be used in routed networks. The *IP loop* mode is not available in *Ethernet endpoint mode*. • *UDP loop*: It works in the same way that the *IP loop* mode but it swaps source and destination UDP ports as well. This mode can be used if there are network devices working at the transport layer like for example firewalls or NAT routers. The *UDP loop* mode is not available in *Ethernet endpoint* mode.

Table 2.7: Local Ethernet / IP Profile

Result	Description
Traffic to loop	Configures which frames are looped in the current test port. The available configurations are: • *All frames*: Loops all frames received in the test port. • *Filtered frames*: Loops frames matching any of the eight receiving filters available for configuration. All non-matching frames are discarded. Use this option if you need a tight control on the looped frames.
Loop broadcast frames	Chooses whether to loop Ethernet broadcast frames. If broadcast loop is disabled, all frames with destination MAC address set to *FF:FF:FF:FF:FF:FF* are discarded. Use this setting if you want to avoid Ethernet broadcast frames to proliferate and potentially flood the network.
Loop ICMP packets	Chooses whether to loop ICMP packets. If ICMP loop is disabled, all ICMP frames (IP protocol number 1) will be discarded. Use this setting to avoid ICMP packet proliferation in the network.

While in loop mode, the test port does not reply ICMP echo request messages and it ignores all other network protocols. The only exception to this rule is ARP. A test port replies ARP requests received from the network. Of course, ARP requests are also looped if the port is configured to do so.

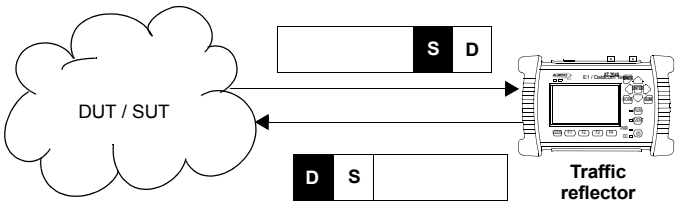

Figure 2.5: This figure illustrates the loopback operation (*MAC loop*). Source and destination addresses are swapped so that they can be processed in the normal way by intermediate switches.

2.5. Testing with an External Clock Reference

Some tests performed by the Ether.Genius / Ether.Sync require a clock reference. Moreover, it is possible that the user is interested in setting the transmission timing source of Port A / Port B transmitters to something different to the default configuration (the internal clock or in the 1000BASE-T slave interface the recovered clock). Clock reference configuration for tests that require an external timing source is carried out through the *Reference clock* menu. Configuration of an external clock reference is not available for Ether.Giga. The source timing candidates are the following ones:

- *Internal*: The test unit clock reference is configured to use the timing from an internal oscillator. This oscillator is a temperature-controlled crystal oscillator (TCXO) which provides a frequency accuracy better than ±2.0 ppm. For Ether.Sync and Ether.Genius the user can optionally replace the TCXO by an oven-controlled crystal oscillator (OCXO) which provides a frequency accuracy about one order of magnitude better than the TCXO.
- *2048 kHz (Port C)*: The clock reference is derived from an external ITU-T G.703 2048 kHz signal. This signal is received through the Port C. For this reason, this option is available only in Ethernet / IP operation modes (*Ethernet endpoint*, *IP endpoint*, *IP through*, *Ethernet cable test)*. Ether.Genius users must use *2048 kHz (DTE port)* for E1 operation modes.
- *E1 (Port C)*: The clock reference is derived from an external ITU-T G.703 2048 kb/s signal. This signal is received through the Port C. For this reason, this option is available only in Ethernet / IP operation modes (*Ethernet endpoint*, *IP endpoint*, *IP through*, *Ethernet cable test)*. Ether.Genius users must use *E1 (DTE port)* for E1 operation modes.
- *1544 kHz (Port C)*: The clock reference is derived from an external 1544 kHz signal. This signal is received through the Port C. For this reason, this option is available only in Ethernet / IP operation modes (*Ethernet endpoint*, *IP endpoint*, *IP through*, *Ethernet cable test)*. Ether.Genius users must use *1544 kHz (DTE port)* for E1 operation modes.
- *T1 (Port C)*: The clock reference is derived from an external ITU-T G.703 1544 kb/s signal. This signal is received through the Port C. For this reason, this option is available only in Ethernet / IP operation modes (*Ethernet endpoint*, *IP endpoint*, *IP through*, *Ethernet cable test)*. Ether.Genius users must use *T1 (DTE port)* for E1 operation modes.
- *10 MHz (Port C)*: This is a 10 MHz unipolar clock, square or senoidal signal received through Port C. This input is not available for the E1 operation modes (Ether.Genius). Ether.Genius users willing to use a 10 MHz input for the E1 operation modes (*E1 endpoint*, *E1 monitor*, *E1 through*) should use the *10 MHz (DTE port)* input instead.
- *PPS (DTE port)*: The clock reference is derived from an external 1 pps signal received through the DTE port with the help of the AT-91 synchronization adaptor. The 1 PPS input can be used both for Ethernet / IP operation modes (*Ethernet end-*

point, *IP endpoint*, *IP through*, *Ethernet cable test*) and Ether.Genius E1 tests (*E1 endpoint*, *E1 monitor*, *E1 through*).

- *2048 kHz (DTE port)*: The clock reference is derived from an external ITU-T G.703 2048 kHz signal received through the DTE port with the help of the AT-91 synchronization adaptor. This external input is available only for the Ether.Genius E1 operation modes (*E1 endpoint*, *E1 monitor*, *E1 through*).
- *E1 (DTE port)*: The clock reference is derived from an external ITU-T G.703 2048 kb/s signal received through the DTE port with the help of the AT-91 synchronization adaptor. This external input is available only for the Ether.Genius E1 operation modes (*E1 endpoint*, *E1 monitor*, *E1 through*).
- *1544 kHz (DTE port)*: The clock reference is derived from an external 1544 kHz signal received through the DTE port with the help of the AT-91 synchronization adaptor. This external input is available only for the Ether.Genius E1 operation modes (*E1 endpoint*, *E1 monitor*, *E1 through*).
- *T1 (DTE port)*: The clock reference is derived from an external ITU-T G.703 1544 kb/s signal received through the DTE port with the help of the AT-91 synchronization adaptor. This external input is available only for the Ether.Genius E1 operation modes (*E1 endpoint*, *E1 monitor*, *E1 through*).
- *10 MHz (DTE port)*: This is a 10 MHz unipolar clock, square or senoidal signal received through the DTE port with the help of the AT-91 synchronization adaptor. This input is available for the Ether.Genius E1 operation modes (*E1 endpoint*, *E1 monitor*, *E1 through*) only.
- *Ethernet (Port A) / Ethernet (Port B)*: A Synchronous Ethernet signal recovered from Port A (or Port B) is used as the external clock reference. The equipment may fail to synchronize to a conventional (asynchronous) Ethernet input depending on the specific frequency offset of the input compared with the free running frequency of the internal clock. For the particular case of a 1000BASE-T input, the interface must be forced to an slave role before the input can be used as a synchronization reference (See section 2.2.4).

The procedure to configure the clock reference input in your Ether.Genius / Ether.Sync unit is detailed in the following steps.

Table 2.8: External Clock Reference Settings

Result	Description
Port C connector	Selects the Port C connector to be used as the clock input / output. The available options are: • *Unbalanced*: BNC unbalanced connector, nominal impedance of 75 Ω (when not configured in high impedance mode). • *Balanced*: RJ-45 balanced connector, nominal impedance of 120 Ω (when not configured in high impedance mode).

Table 2.8: External Clock Reference Settings

Result	Description
Port C termination	It configures the input / output impedance of the clock interface where the *Port C* is going to be connected. The available configurations for this field are: • *Endpoint*: This connection represents a network termination point with the nominal impedance (75 Ω for the unbalanced port and 120 Ω for the balanced one). The expected attenuation is the theoretical cable attenuation which increases with the frequency square root. • *-20 dB monitor*: 20 dB protected monitoring point. This is a connection point that is isolated from the network and it is specially suited for monitoring purposes. A flat attenuation of 20 dB is expected for these points.
DTE port termination	It configures the input / output impedance of the clock interface where the DTE port is going to be connected through the AT-91 synchronization adaptor. The available configurations for this field are the same that for the *Port C termination* setting.

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Go to *Reference clock*.
3. Configure *Input clock* to one of *Internal*, *2048 kHz (Port C)*, *E1 (Port C)*, *1544 kHz (Port C)*, *T1 (Port C)*, *10 MHz (Port C)*, *PPS (DTE port)*, *2048 kHz (DTE port)*, *E1 (DTE port)*, *1544 kHz (DTE port)*, *T1 (DTE port)*, *10 MHz (DTE port)*, *Ethernet (Port A)*, *Ethernet (Port B)*.
4. If you have chosen one of *2048 kHz (Port C)*, *E1 (Port C)*, *1544 kHz (Port C)*, *T1 (Port C)* or *10 MHz (Port C)* in the previous step, configure *Port C connector* and *Port C termination*. If you have chosen one of *2048 kHz (DTE port)*, *E1 (DTE port)*, *1544 kHz (DTE port)*, *T1 (DTE port)* or *10 MHz (DTE port)* in the previous step, configure *DTE port termination*.

2.6. Using the Clock Reference Output

The Ether.Genius / Ether.Sync clock reference output enables you to synchronize any external equipment with a clock signal generated by the tester or to connect the internal /recovered clock to an oscilloscope, spectrum analyser or other equipment. The reference clock output is therefore very useful for many tests related with network synchronization. These are the clock reference output ports and formats allowed by Ether.Genius and Ether.Sync:

- *2048 kHz (Port C)*: The clock reference output is encoded as a ITU-T G.703 2048 kHz signal. This signal is transmitted through the Port C. For this reason, this option is available only in Ethernet / IP operation modes (*Ethernet endpoint*, *IP endpoint*, *IP through*, *Ethernet cable test)*. Ether.Genius users must use *2048 kHz (DTE port)* for E1 operation modes.
- *PPS (DTE port)*: The clock reference output is encoded as a 1 pps signal transmitted through the DTE port with the help of the AT-91 synchronization adaptor. The 1 pps output can be used both for Ethernet / IP operation modes (*Ethernet endpoint*, *IP endpoint*, *IP through*, *Ethernet cable test*) and Ether.Genius E1 tests (*E1 endpoint*, *E1 monitor*, *E1 through*).

The configuration procedure to enable the reference clock output is detailed in the following steps:

1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Go to *Reference clock*.
3. Configure *Output clock* to either *2048 kHz (Port C)* or *PPS (DTE port)*.
4. If you have chosen *2048 kHz (Port C)* in the previous step, configure *Port C connector* and *Port C termination*.

Chapter 3
Cable Tests

The Albedo Telecom Ether.Genius / Ether.Sync / Ether.Giga testers have the ability to check wiring and performance of Ethernet cables or optical fibres to make sure they will operate as expected when connected to the network.

Cable testing can be used to estimate the cable length, recognise cabling types or detect wiring faults. Performance of optical fibers is measured in terms of the received optical power.

3.1. Electrical Test

Electrical cable test requires an special operation mode which prevents the unit sending any frame (See section 2.1). When the unit is configured in *Ethernet cable test*, mode, the port mode becomes *Cable test* (port A) and *Link* (port B) and there is no way to change the port mode unless the general operation mode is set to something different of *Ethernet cable test*.

In this operation mode the Ether.Genius / Ether.Sync / Ether.Giga unit could have a network link in the remote end of the device under test (the cable) or not. Depending on how the cable is connected, it would be possible to compute different sets of results.

3.1.1. Basic Principles of Ethernet Cable Wiring

Twisted pair cables for Ethernet LAN applications generally come in groups of four pairs (8 wires). Only two of the four pairs carry information at 10 and 100 Mb/s but all four pairs are simultaneously used in 1 Gb/s links. There are many possible pair interconnections that would work, but only two of them are standard. These are known as T-568A and T-568B wire maps. Usually, all four pairs are always connected, but for 10 and 100 Mb/s operation only pairs 2 and 3 are used. 1 Gb/s port and cable wirings are designed to be backwards compatible with slower 10 Mb/s and 100 Mb/s interfaces so that the same gigabit cable can be used for all bit rates. Some cables designed to operate at 10 Mb/s and 100 Mb/s may have only connections necessary for transmission at these speeds. These cables are not compatible with the 1 Gb/s bit rate.

Ethernet cables may have poor performance or may not work at all if wires are not properly connected.

All the different Ether.Genius / Ether.Sync / Ether.Giga test results are provided for each *Media Dependent Interface* (MDI). An MDI corresponds with a single Ethernet pair transceiver. There are four of them in 1000BASE-T (MDI-0, MDI-1, MDI-2 and MDI-3) and two in 100BASE-TX / 10BASE-T (MDI-0 and MDI-1). Using the MDIs to supply information about wiring is the most logical choice: Pin assignment in a connector may change but MDIs are always the same and at logical level, communication between MDIs is very simple: MDI-0 is connected to MDI-0, MDI-1 is connected with MDI-1, etc.

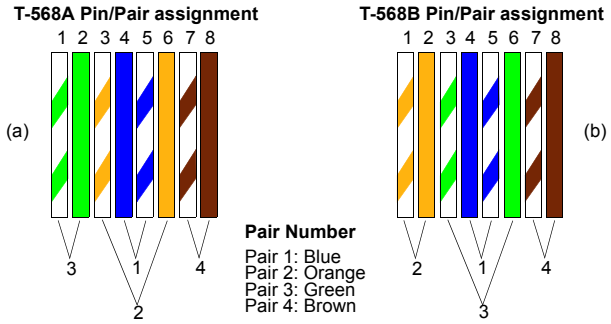

Figure 3.1: T-568A and T-568B wiring standards. Pair numbering (1, 2, 3, 4) is a way to identify pairs in the cable and is not to be confused with MDI numbering (MDI-0, MDI-1, MDI-2, MDI-3).

Old Ethernet stations had a fixed pinout for MDI-0, MDI-1, MDI-2 and MDI-3. This fixed assignment for stations corresponds to the *MDI* status. These older stations were connected to Ethernet hubs and switches through a straight cable and to a second station through a crossover cable. For this reason, MDI pair assignment in hubs and switches was required to be complementary to the wiring of stations. This second pair assignment corresponds to the *MDIX* status:

Table 3.1: Pair assignment for MDI and MDIX port modes

	MDI-0	MDI-1	MDI-2	MDI-3
MDI	1-2	3-6	4-5	7-8
MDIX	3-6	1-2	7-8	4-5

In old devices with static MDI pair assignment, the description of the interface could be based either on MDI-0 to MDI-3 or pair 1 to pair 8 but with most of the current ethernet

devices the situation is different. Modern Ethernet stations and switches can switch their ports between *MDI* and *MDIX* status before link establishment to make sure that an MDI port is connected with a remote MDIX port through an straight cable and that an MDI (or MDIX) port is connected to a second MDI (or MDIX) port through a crossover cable.

The advantage of this procedure is that there is no need to think about which cable (crossover or straight) is required for interconnection of network equipment and stations. However, if you need to know the way the cable is wired, then it is mandatory to know the MDI / MDIX status of the near and far ends.

Table 3.2: Crossover status (*MDI* / *MDIX*) and cable Wiring

Near end	Far end	Cable
MDI	MDI	Crossover
MDI	MDIX	Straight
MDIX	MDI	Straight
MDIX	MDIX	Crossover

3.1.2. Connecting the Tester

Cable testing with the *Ethernet cable test* mode is different to the other modes because in this case the test equipment does not generate any digital framed or unframed signal. It relies on the remote end to supply some information about the wiring and other metrics related with the physical transmission medium. The amount of information to be supplied depends on the communications interface. It is not the same for Gigabit Ethernet than for Fast Ethernet interfaces. If there is no link from the remote end, the equipment attempts to find out the reason (short circuit, open circuit...).

The test cable is available for Port A only. The Port B could be used as an auxiliary port to determine whether a cable is straight or crossover. Cable tests can only be run from the native RJ-45 interfaces, they are not available through electric SFP modules.

The best way to run the cable test is to connect the cable under test between port A and port B in Ether.Genius / Ether.Sync / Ether.Giga. You can use other test configurations like a connection between Port A and a switch but in this case you will be unable to get all the information from the cable. Here you have some details about the the available connection modes for the cable test and the results obtained from each of them:

- Port A - Port B closed loop: *Fault*, *Crosstalk*, *Distance (m.)*, *Crossover*, *Polarity*, *Skew (ns)*, *Wiring*.
- Port A connection with remote link: *Fault*, *Crosstalk*, *Distance (m.)*, *Crossover*, *Polarity*, *Skew (ns)*.
- Port A connection without link: *Fault*, *Crosstalk*, *Distance (m.)*.

In case a fault is found in some pair (*Open*, *Short*), some results may not be available for the corresponding MDI. If no fault is found then the *Distance (m)* field is empty. The *Skew (ns.)* test result is displayed only in 1000BASE-T interfaces.

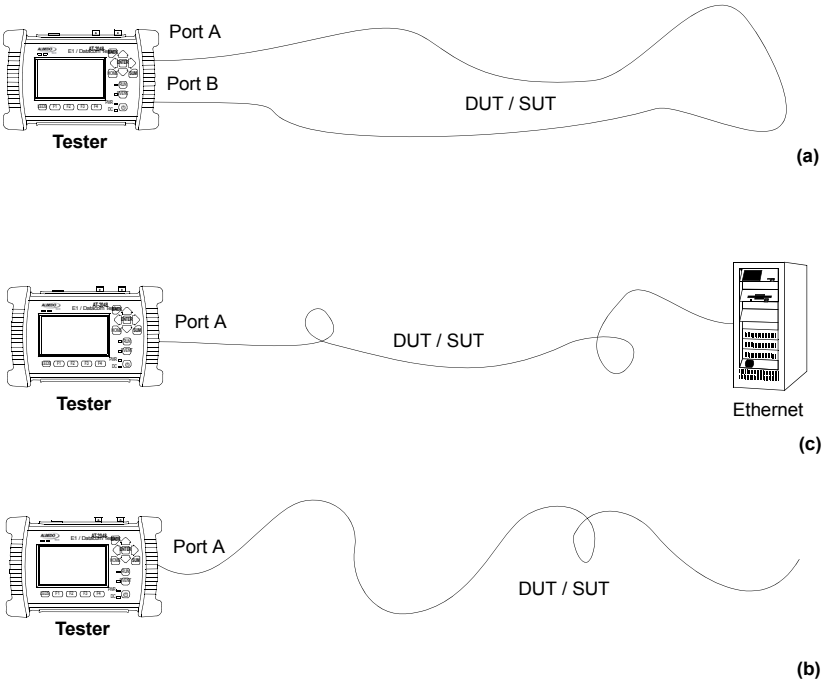

Figure 3.2: Basic connection setup for Ether.Genius, Ether.Sync and Ether.Giga testers: (a) Port A - Port B closed loop, (b) Port A connection with remote link, (c) Port A connection without link.

3.1.3. Running the Test

Once the tester has been connected to the network in any of the supported test configurations. It is required to configure and run the cable test by means the tester user interface. It is assumed that the *Connector* setting is configured to *Electrical* in the test ports (See section 2.2.1). The test procedure is as follows:

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Select *Mode* to enter in the mode selection menu

3. Choose *Ethernet cable test*. Confirm by pressing ENTER.
4. From the *Home* panel, go to *Results*,
   The port setup panel is displayed.
5. Select *Port A* to enter in the test Port A specific results panel.

Table 3.3: Common cable wiring problems

Description	Diagram	Diagnostic
Inverted cable. Polarity is inverted in both ends of the same pair	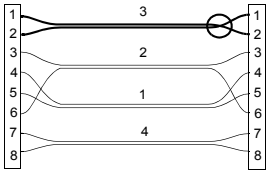	The *Polarity* result for de corresponding MDI (MDI-0, MDI-1, MDI-2 or MDI-3) is *Negative* rather than *Positive.*
Pairs 1 and 2 are wired to the wrong pins in the connector	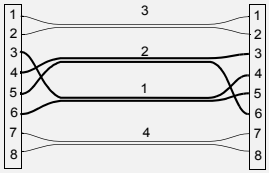	The *Fault* result for MDIs corresponding to pairs 1 and 2 is *Open*. *Distance (m)* displays the distance to the far end. The *Crosstalk* shows that pairs 1 and 2 are coupled.
There is a short circuit between the conductors in pair 2	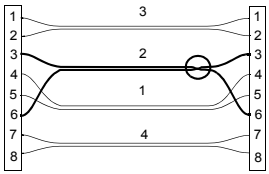	The *Fault* result for the MDI corresponding to pair 2 displays *Short*. *Distance (m)* displays the distance to the short circuit.
One of the cables of port A is broken and it contains an open circuit	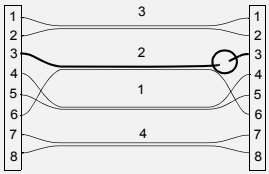	The *Fault* result for the MDI corresponding to pair 2 displays *Open*. *Distance (m)* displays the distance to the open circuit.

Table 3.3: Common cable wiring problems

Description	Diagram	Diagnostic
Swapped pair, connections are OK but wires of the same connection are twisted in different pairs.	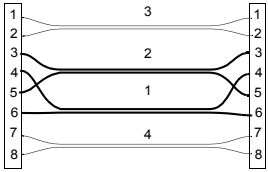	No fault is detected in any pair (*Fault* displays *OK* for all MDIs) but if the cable is long enough it will result in significative crosstalk between coupled pairs.
Miswired cable	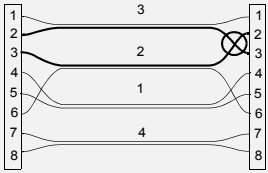	The diagnostic of the cable shows crosstalk between the MDIs corresponding to pairs 2 and 3.

6. Go to Cable
   The Ethernet cable test result panel is displayed
7. Run the test by pressing the *RUN* test and wait a few seconds to the end of the measurement.
8. Check the *Fault*, *Crosstalk*, *Distance (m.)*, *Crossover*, *Polarity*, *Skew (ns.)* and *Wiring* test results for each MDI-*n*.

Table 3.4: Cable test results

Result	Description
Fault	Displays information about faults found in the corresponding MDI. In case a fault is found the *Fault* indication could be either *Open* or *Short*. • *Open* means that it has been found an open circuit in the remote end. An open circuit is declared when the impedance in the remote end is very large (the reflection coefficient is close to 1). • *Short* means that it has been found a short circuit in the remote end. The short circuit is declared if the impedance in the remote end is zero or close to zero (the reflection coefficient is close to -1).

Table 3.4: Cable test results

Result	Description
Crosstalk	Crosstalk is detected when there is electromagnetic coupling between pairs. Coupling could be inductive / capacitive if there is a twisting defect in the pair or it can be resistive if there is a problem with the dielectric medium between conductors. Some wiring faults between the cable and the connector may cause crosstalk too. The crosstalk is indicated in the equipment as a collection of numbers separated by '-' for each MDI. For example if MDI-0 and MDI-3 are coupled, then the crosstalk result for MDI-0 and MDI-3 will be 0-3.
Distance (m.)	The *Distance (m.)* result displays the distance to an open circuit / short circuit fault in with an accuracy of ±1 m. If you there is no fault, no distance is displayed. The maximum range of the distance test is 100 m. To measure the cable length, just leave the far end disconnected and run the test. Ether.Genius / Ether.Sync / Ether.Giga will detect an open circuit in all MDIs and the *Distance (m.)* results displays the cable length.
Crossover	Displays whether the local *MDI-n* is in straight (MDI) or crossover (MDIX) status. Note that *Crossover* is not really a cable result, it is the local port MDI/MDIX status. This status may be random and depends on the cable and the remote port. Therefore this result may be modified if the cable is disconnected and reconnected again.
Polarity	Polarity could be positive or negative for each *MDI-n.* A negative polarity indicates that the pair connects pins of inverted polarity in the local and remote end. Positive polarity means that local and remote pins have the same assigned polarity.
Skew (ns.)	Relative propagation delay, expressed in nanoseconds, experienced by the pair associated to the MDI and compared to the MDI that has minimum propagation delay. That means that the Skew (ns.) is always zero for at least one MDI. The skew result is not measured for 10 Mb/s and 100 Mb/s interfaces.

Table 3.4: Cable test results

Result	Description
Wiring	A cable could be designed to cross at the remote end the pairs used for transmitting and receiving (crossover cable) or not (straight cable). The wiring result checks whether a cable is straight or crossover. This test result depends on the local and remote MDI / MDIX status and it can only be determined if the cable is connected through test port A and B.

3.2. Measuring Optical Power

Ether.Genius / Ether.Sync / Ether.Giga report the transmitted and received optical power if they are equipped with SFPs supporting this measurement (See section 2.2.2). The following description assumes that the *Connector* setting is configured to *Optical* in the test ports (See section 2.2.1). The test procedure is as follows:

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Select *Mode* to enter in the mode selection menu
3. Choose *Ethernet endpoint, IP endpoint or IP through*. Confirm by pressing ENTER.
4. From the *Home* panel, go to *Results*,
   The port setup panel is displayed.
5. Select either *Port A or Port B* to enter in the port specific results menu.
6. Go to *SFP information*.
7. Check the results of *TX optical power* and *RX optical power*.

Chapter 4
Traffic Generation

One of the key features of Ether.Genius / Ether.Sync / Ether.Giga tester families is the ability to generate traffic with deterministic and random bandwidth profiles. The traffic generation feature can be used to stress the network, simulate user traffic and, if a suitable payload is configured, to measure critical network performance parameters like bit errors, packet loss or latency.

Ether.Genius / Ether.Sync / Ether.Giga test equipments have eight independent full featured traffic generators attached to the main test port (Port A). Each traffic flow may be configured with specific encapsulation and addressing parameters thus providing great versatility in all applications requiring Ethernet and IP traffic generation.

4.1. Generation of Ethernet Traffic

In Ether.Genius / Ether.Sync / Ether.Giga testers, generation of custom Ethernet frames is available for Port A in *Ethernet endpoint* mode through the *Frame*, *Bandwidth profile* and *Payload* settings for each of the eight available traffic flows. This is a short description of the Ethernet traffic generation menus:

- *Frame*: Configures the encapsulation and MAC addresses. If the Ethernet frames have any VLAN tag, this menu configures the VID and priority for these tags.
- *Bandwidth profile*: Sets the traffic generation statistics. There are four different generation profiles to choose: *Constant*, *Periodic burst*, *Ramp* and *Random*.
- *Payload*: This menu is used to set the payload to be inserted in the generated Ethernet frames. The SLA payload enables the user to measure delay, jitter and packet loss. The BERT payload (flow 1 only) is used for BER testing in framed interfaces.

Frame generation capability in Ether.Genius / Ether.Sync / Ether.Giga is controlled by RUN button. That means that no test traffic is generated if you don’t press RUN. However, the tester may generate signalling traffic or reply to certain messages like ARP or ICMP echo requests / replies even if there is not an ongoing test. Some automatic tests like the RFC 2544 or the eSAM have their own internal traffic generation dynamics but they are controlled by the RUN button as well.

4.1.1. Physical Layer Settings

Before starting any frame generation test, the equipment must be connected to the network and the electrical and optical physical layer must be correctly configured. Ethernet technology has been designed to keep physical layer configuration to the minimum. But there are at least two settings you may need to check before you get a link from the DUT / SUT. These settings are the *Connector* (See section 2.2.1) and the *Auto-negotiation* (See section 2.2.4). You will know that Port A is prepared for traffic generation and analysis when the *1000, 100* or *10* LED (See section 5.3) is displayed in green colour.

4.1.2. Frame Settings

Most of the Ethernet frame fields are available for configuration in Ether.Genius / Ether.Sync / Ether.Giga testers. Before configuring these fields it is necessary to tell to the tester which frame structure is going to be used for traffic generation. The *Frame type* is a port-wide setting. Once you choose an specific framing for your traffic, all streams you define for the port carry the same framing structure. The *available Frame type* settings are:

Table 4.1: Ethernet Frame Settings

Setting	Description
Encapsulation	This field configures the way the data is encapsulated in Ethernet frames for transmission in the current stream. The allowed encapsulations are the following ones: • *None*: A DIX or IEEE 802.3 frame carries the test data, depending on the current value of the *Frame type* setting. • *VLAN*: Transmitted frames are labelled with an IEEE 802.1Q frame tag. Settings related with configuration of the VLAN tag are enabled when this option is selected. • *Q-in-Q*: Transmitted frames carry two VLAN tags, one service provider tag (S-VLAN) and a customer tag (C-VLAN). The C-VLAN is identified by the normal IEEE 802.1Q Ethertype and the S-VLAN carries one of the not-standard Ethertypes. • *IEEE 802.1ad*: Transmitted frames carry two VLAN tags. It is similar to the Q-in-Q encapsulation but this option follows strictly the standard IEEE 802.1ad encapsulation for Provider Bridges (PB). Specifically, the IEEE 802.1ad carries the special 0x88a8 Ethertype within the S-VLAN.

Table 4.1: Ethernet Frame Settings

Setting	Description
Source MAC address from	Establishes the origin of the source MAC address for the current stream. There are two possible settings: • *Local*: The source address is set to the factory MAC address assigned to the port. Use this setting if there is no other requirement. • *Manual*: The source address is set to the value configured in *Source MAC address*. Use manual MAC addresses if you want to simulate traffic generated by an equipment different to the tester or, in multi-stream operation, to simulate traffic transmitted from different stations. Most of the times you will want to avoid duplicated addresses in your network. For this reason, make sure that no other equipment is using the manually configured MAC address.
Source MAC address	Source MAC address carried by the frames generated in the current stream if *Source MAC address type* is set to *Manual*. Anything from 00:00:00:00:00:00 to ff:ff:ff:ff:ff:ff is allowed.
Des. MAC address from	Establishes the origin of the destination MAC address for the current stream. There are three different settings available for configuration: • *ARP*: Uses the Address Resolution Protocol (IETF RFC 826) to configure the destination MAC address without user intervention. The ARP requires the IPv4 destination address to be previously configured to work. For this reason, ARP is available only in *IP endpoint* mode. • *Manual*: The destination address is set to the value configured in *Destination MAC address*. • *Range*: Test data in the current stream is transmitted to a group of MAC addresses configured with *Destination MAC address* and *Address number within range*. Use this option if you want to deliver the test data sequentially to many different destinations.
Destination MAC address	Destination MAC address carried by the frames generated in the current stream if *Des. MAC address type* is set to *Manual*. If *Des. MAC address type* is set to *Range*, this field contains the first destination MAC address within the range. Anything from 00:00:00:00:00:00 to ff:ff:ff:ff:ff:ff is allowed for this field.

Table 4.1: Ethernet Frame Settings

Setting	Description
Address range size	Configures the number of MAC addresses within an address range. This control is valid only if *Des. MAC address type* is set to *Range*. In this case, the ethernet frames transmitted in the current stream will contain as many destination addresses as previously configured in this field. The destination MAC address is increased by one unit for each transmitted frame starting with the value configured in *Destination MAC address*. If there are no more addresses left in the range, transmission returns to the initial address and starts the process from the beginning.
Ethertype	Ethertype value carried by the frames generated in the current stream. This value is found within the Ethernet *Type* header field in DIX / Ethernet II frames or within the LLC / SNAP header in IEEE 802.3 frames. Depending on the configuration, the *Ethertype* value is fixed and cannot be set by the user. If the operation mode is *IP endpoint*, the Ethertype is automatically configured to 0x0800 (Internet Protocol, version 4). If the payload type is configured to *SLA* in *Ethernet endpoint* mode, the Ethertype is set to 0x8902 (IEEE 802.1ag / ITU-T Y.1731 OAM) to account for the special structure of the Ethernet SLA measurement payload.
C-VID	VLAN identifier assigned to tagged frames (IEEE 802.1Q) or C-VLAN identifier for double tagged frames (IEEE 802.1ad, Q-in-Q). In frames with two VLAN tags, the C-VID usually accounts for the VLAN structure corresponding to the customer network. Any value within 0 to 4096 is allowed for this field.
C-VLAN priority	3-bit class of service (CoS) field defined to set frame groups with different priorities or to provide specific treatments to special frames within a network or an administrative domain. This field is carried by the Q-tag of Ethernet frames with a single tag or by the C-tag of Ethernet frames with two tags. Any value from 0 to 7 is allowed for this field. Specific actions to be carried out on frames with different CoS labels depend on the network and the service provider.

Table 4.1: Ethernet Frame Settings

Setting	Description
S-VLAN TPID	Ethertype to be associated to the S-VLAN tag in Q-in-Q frames. Four different values are possible: 0x8100, 0x9100, 0x9200 and 0x9300. It the encapsulation is set to IEEE 802.1ad, the S-VLAN Ethertype is automatically set to 0x88a8 and this field is not available for configuration.
S-VID	VLAN identifier assigned to the S-tag in double tagged frames (IEEE 802.1ad, Q-in-Q). In frames with two VLAN tags, the S-VID usually accounts for the VLAN structure corresponding to the service provider network. Any value within 0 to 4095 is allowed for this field.
S-VLAN priority	3-bit class of service (CoS) field defined to set frame groups with different priorities or to provide specific treatments to special frames within a network or an administrative domain. This field is carried by the S-tag (service provider tag) of Ethernet frames with two tags. Any value from 0 to 7 is allowed for this field. Specific actions to be carried out on frames with different CoS labels depend on the network and the service provider.
Drop-eligible indicator	This is a single bit field that is used to mark drop eligible frames. These frames are usually dropped first when congestion is detected in a network node. The Drop eligible operator is carried within the S-tag of IEEE 802.1ad frames.
Frame size	Ethernet MAC frame size including the destination MAC address, source MAC address, type / length field, payload, FCS and any VLAN tag carried by the frame. Anything between 64 B and 10000 B is allowed but frames longer than 1518 B (without VLAN tags and MPLS labels) are out of the IEEE 802.3 standard. It is possible to generate frames longer than the port Maximum Transmission Unit (MTU) but these frames are considered oversized frames when they are analysed by the tester. To avoid an *OverS* anomaly in this case, increase the value of the port MTU.

- *DIX*: Port A generates *DEC, Intel, Xerox* (DIX) frames, also known as Ethernet II frames. DIX / Ethernet II frames encode the payload type in the *Type* frame field.

This is the most common framing format found in real networks: For example, RFC 894 mandates a DIX / Ethernet II frame structure with the *Type* field set to 0x0800 for IPv4 encapsulation.

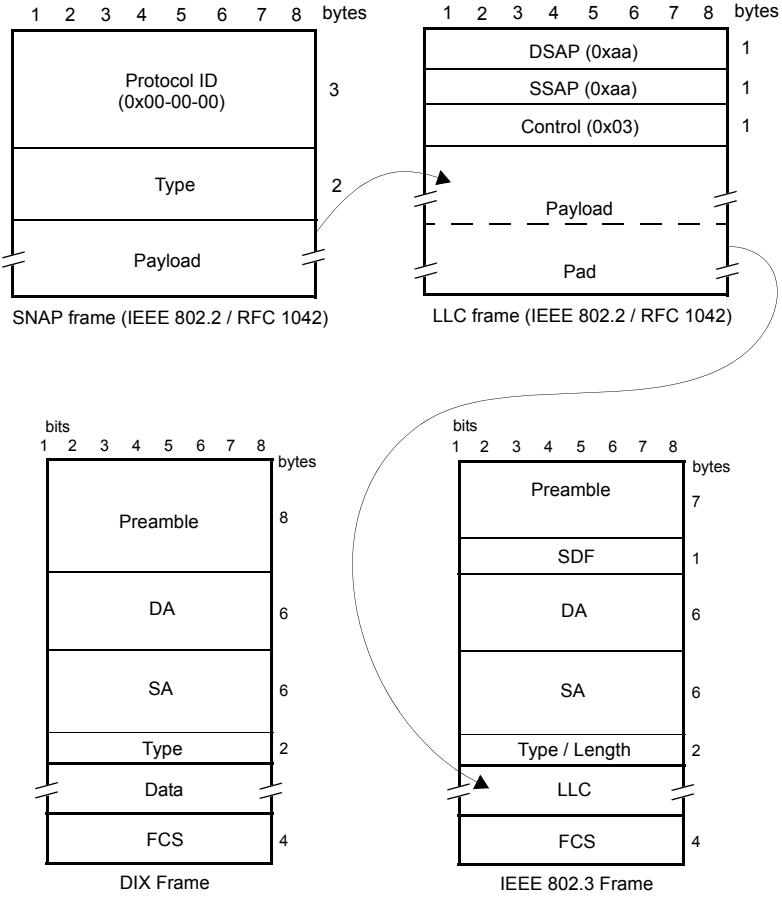

Figure 4.1: MAC frame structure: IEEE 802.3 and DIX

- *IEEE 802.3*: Frame format defined in IEEE 802.3 standard. It is similar to the DIX frame but it specifies a different usage for the Type field that is renamed to Type / Length value. If Type / Field is larger than 0x0600 then it has the same meaning than the DIX / Ethernet II field but otherwise it specifies the frame length in bytes.

IEEE 802.3 frames leave to the IEEE 802.2 *Logical Link Control* (LLC) the specification of the payload type.

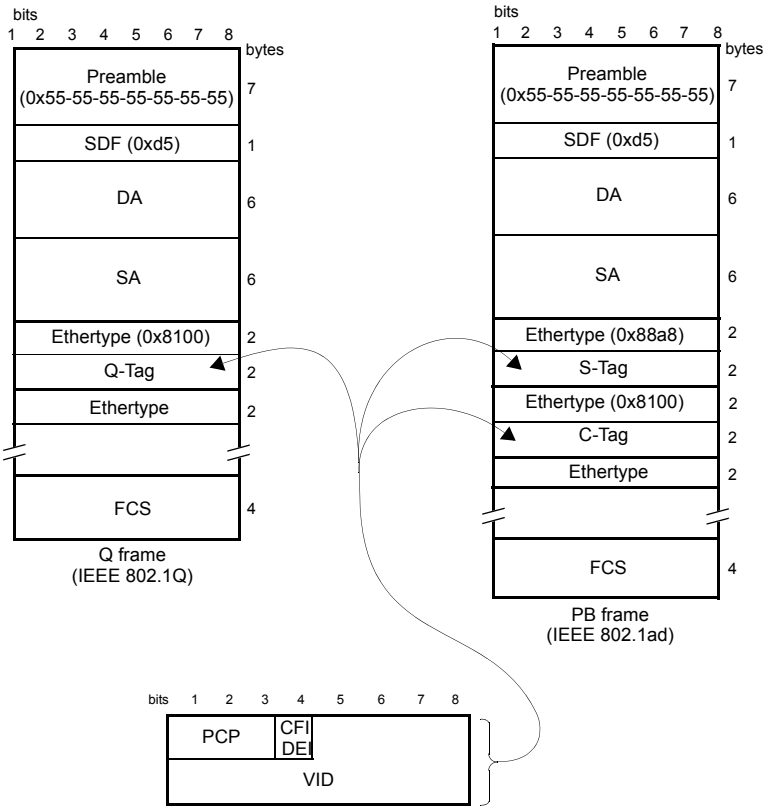

Figure 4.2: IEEE 802.1Q y IEEE 820.1ad frame structures.

The second port-wide setting to be configured is the Maximum Transmission Unit (MTU). This setting is relevant for the analyser only and it configures the largest frame size accepted without declaring the *OverS* defect. Standard IEEE 802.3 specifies an MTU of 1518 bytes for ordinary Ethernet frames but 1522 is admitted for VLAN frames and 1526 valid for frames carrying two VLAN tags (IEEE 802.3ad, Q-in-Q). Some switches provide support for much larger frames known as jumbo frames. These frames are more efficient because the ratio of header bytes to payload bytes is smaller for larger frames but they are currently not accepted by any international standard.

The following steps illustrate the frame configuration procedure in any of the Ether.Genius / Ether.Sync / Ether.Giga testers. Both the port-wide and flow-specific configuration is included.

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
3. Select either Port A or Port B to enter in the port specific configuration.
   Note: Most of the frame configuration settings are not available from Port B because traffic generation is not available from this port.
4. Enter in the *Frame* menu.
   All settings related with frame configuration are displayed.
5. Configure the correct MTU with the help of the *MTU* menu. You may want to set the MTU to 1518 bytes for traffic analysis in line with IEEE 802.3 or to other value to allow jumbo frames. The maximum allowed MTU is 10,000 bytes.
6. Select one of the traffic flows between *Flow #1* and *Flow #8* to enter in the flow specific configuration.
7. Configure the encapsulation you are going to use in the generated frames. Basically, the *Encapsulation* menu sets the number of VLAN tags to be included in the generated frames.
8. Enter the source MAC address with the help of the *Source MAC address from* and *Source MAC address* controls. You can configure the factory MAC address as the source address for the generated frames or enter a custom address.
9. Enter the destination MAC address or addresses by using the *Destination MAC address from*, *Destination MAC address* and *Address range size*. If you choose to generate a destination address range you will be requested to enter the number of addresses that made up the range.
10. Configure the *Ethertype* value.
    Note: Some frame structures require an specific value of the Ethertype. This field cannot be configured in this case.
11. If you are using frames carrying one (IEEE 802.1Q) or two (IEEE 802.1ad, Q-in-Q) VLAN tags, enter the *C-VID* and *C-VLAN priority*.
12. If you are using frames carrying two VLAN tags (IEEE 802.1ad, Q-in-Q), enter the *S-VID*, *S-VLAN priority* and *Drop-Eligible Indicator*.
13. If you are generating not-standard Q-in-Q frames, set the *S-VLAN TPID* to one of the allowed values.
14. Configure the frame length to the correct value with the help of *Frame size*.
15. If necessary, repeat the specific flow configuration for one or more traffic flows (*Flow #1* to *Flow #8*) from the *Frame* menu.

4.1.3. Configuring the Bandwith Profile

In the same way that the *Frame* menu configures the frame format for each of the available traffic flows, the Bandwidth profile sets how many frames are transmitted and how transmission events are distributed in time. The simplest is to generate frames with a constant bit rate specified in frames per second, bits per second or as a percentage of the total transmission channel capacity. However, Ether.Genius / Ether.Sync / Ether.Giga provide other alternatives to the constant transmission like the periodic burst and ramp transmission or random transmission with Poisson statistics.

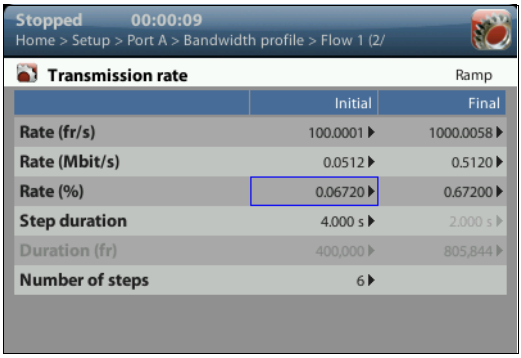

Figure 4.3: Albedo Ether.Genius / Ether.Sync / Ether.Giga bandwidth profile configuration panel.

The bandwidth profile settings are available only in port Port A because the traffic generator is not available in Port B. The procedure to configure the bandwidth profile in a traffic flow is as follows:

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
3. Select Port A to enter in the port specific configuration.
   Note: There is no bandwidth profile configuration for Port B because Port B is unable to generate synthetic traffic.
4. Enter in the *Bandwidth profile* menu.
5. Select one of the traffic flows between *Flow #1* and *Flow #8* to enter in the flow specific configuration.
   All configuration items related with the bandwidth profile are displayed.
6. Configure the transmission mode to one of the available profiles with the help of the *Mode* control.
7. Configure the transmission rate parameters with *Transmission Rate*.
   Note: Depending on the current transmission mode you will be requested to enter

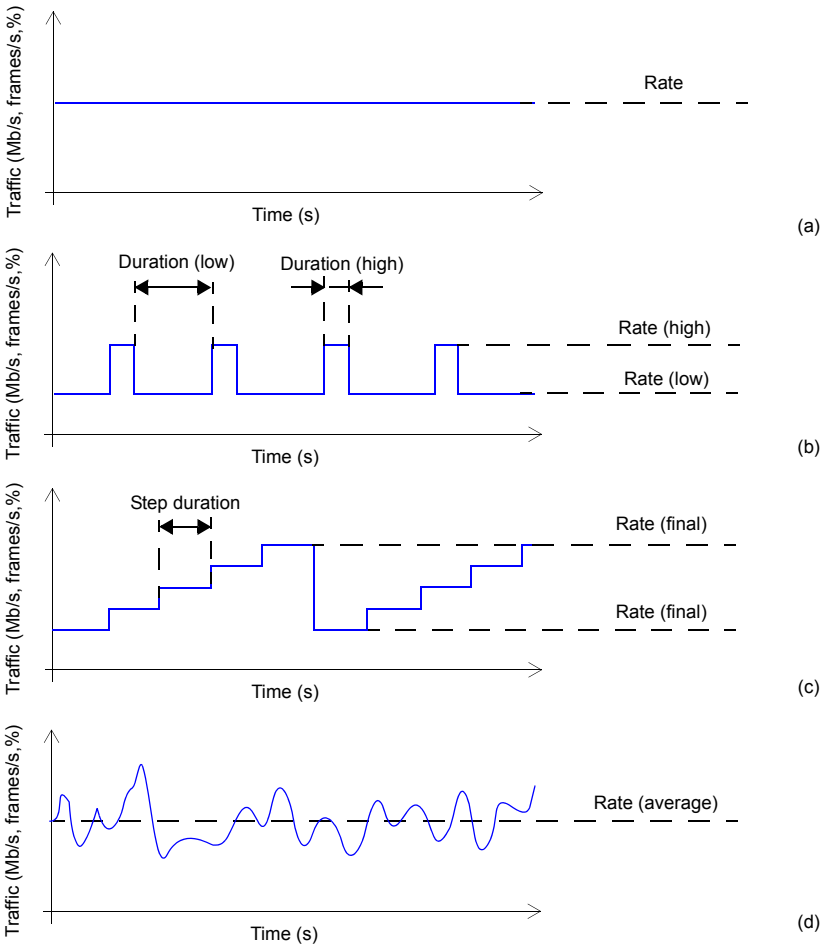

Figure 4.4: Bandwidth profiles for Albedo Ether.Giga / Ether.Genius / Ether.Sync testers: (a) Continuous traffic generation, (b) Periodic burst generation, (c) Ramp generation, (d) Random traffic generation with Poisson probability distribution.

different traffic parameters in the *Transmission Rate* panel.

Note: Changing some transmission parameter may affect the value of other parameters previously configured in the same panel. For example, setting the transmission rate in frames per second modifies the rate in bits per second and the percentage value of the transmission rate.

Note: If the channel capacity varies, the transmission rates configured as percent-

ages of the overall channel capacity are kept to se same value but the bits per second and frames per second are recomputed for the new channel capacity.

Table 4.2: Ethernet Payload Settings

Setting	Description
Mode	Configures the traffic shape to be used by the traffic generator in the current stream. There are five possible generation modes for this for the bandwidth profile • *Off*: No frames are transmitted in the current stream. Use this setting if you want to disable traffic generation in the current stream but you don't want to globally disable generation in the test port. • *Continuous*: Frames is transmitted at a constant speed to match a value configured in bits per second, frames per second or a percentage of the line capacity. • *Periodic burst*: Traffic generation is distributed in periodic bursts of fixed length. Between traffic bursts the user may choose to generate background traffic or disable traffic generation. • *Ramp*: Generates traffic that increases its bit rate with time in steps. The number of steps and step duration are configured by the user. Minimum and maximum traffic generated in the ramp are user configurable as well. Ramp generation is periodic. Traffic generator is restarted when it finishes with the last step of an specific ramp. • *Random*: The number of frames generated per time unit is a Poisson random variable. This is equivalent to say that the distance between two consecutively generated frames is an exponential random variable. Use the Random generation profile to generate traffic that resembles network traffic as much as possible.

Table 4.2: Ethernet Payload Settings

Setting	Description
Transmission rate	This control displays an editable table that enables the user to enter the parameters associated with the traffic to be generated by the current stream. Parameters to be configured depend on the current bandwith profile generation mode: • *Continuous* traffic: The relevant bandwidth parameter is the transmission *Rate* configured in *fr/s*, *Mb/s* or *%*. • *Periodic burst*: Values to be entered are the high and low transmission rates (in *fr/s*, *Mb/s* or *%*) and the high and low durations expressed in seconds or frames. • *Ramp*: Relevant configuration parameters are the initial and final transmission rates (in *fr/s*, *Mb/s* or *%*), the *Step duration* configured in seconds and the *Number of steps*. • *Random*: The bandwidth parameter to be configured is the average transmission rate in *fr/s*, *Mb/s* or *%*.

8. If necessary, repeat the bandwidth profile configuration process for one or more traffic flows (*Flow #1* to *Flow #8*) available from the *Bandwidth profile* menu.

Test traffic generation does not start immediately after setting the bandwidth profile parameters. Traffic generation requires a test to be started with the RUN button.

4.1.4. Choosing the Test Payload for Ethernet

Traffic generated by Ether.Genius / Ether.Sync / Ether.Giga is synthetic. It does not contain any real user data. In fact, the user payload of the internally generated frames is replaced by a test payload. Many times, test payloads are much more than dummy bit sequences designed to replace the user traffic. Test payloads may contain time stamps or sequence numbers that determine which test metrics are available from the result panels or which tests will be run. For this reason, configuration of the right test payload is important to get the required results.

Selection of the test pattern is relevant both for the generator and the analyser. When you generate a test payload or pattern in Port A, the same port is automatically configured so that it is waiting for frames carrying the same pattern in the receiver. Settings related with test payload / pattern selection are available both in Port A and Port B. The procedure to select the test payload in the tester is as follows:

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
3. Select either *Port A* or *Port B* to enter in the port specific configuration.
4. Enter in the *Payload* menu.

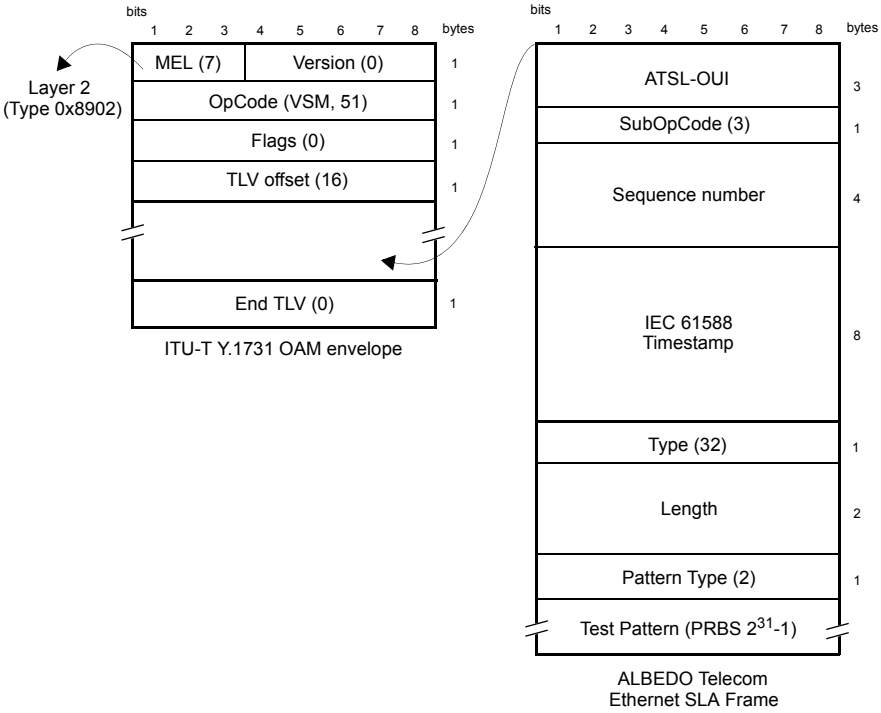

Figure 4.5: Albedo payload for SLA tests (Ethernet Endpoint mode).

5. Select one of the traffic flows between *Flow 1* and *Flow 8* to enter in the flow specific configuration.
   All settings related with payload configuration in the current flow are displayed.
6. Choose one of *BERT*, *SLA* or *All zeroes* in *Payload type*.
   Note: BER is available only for flow 1.
7. If you have configured *Payload type* to BER, choose the bit pattern you are going to use for generation (Port A) and analysis (Port A and Port B) with the help of the *BERT patterns* control.
8. If you have configured *BERT patterns* to *User*, enter a 32-bit test pattern in *User payload* in hexadecimal format.

9. If necessary, repeat the payload configuration process for one or more traffic flows (*Flow #1* to *Flow #8*) available from the *Payload* menu.

Table 4.3: Ethernet Frame Settings

Setting	Description
Payload Type	• *BERT:* The payload content is set to a bit pattern suitable for measuring the Bit Error Ratio (BER). The tester includes support for two different kinds of BERT pattern: Pseudo-Random Bit Sequences (PRBSs) or 32-bit user configurable patterns. BERT generation and analysis over framed interfaces is supported by flow 1 only. • *SLA:* This is the payload to be used to measure latency, packet loss and all the SLA metrics derived from them. If the current operation mode is set to *Ethernet endpoint*, the SLA payload constitutes a proprietary extension of the Operation, Administration and Maintenance (OAM) protocol for Ethernet defined in ITU-T Y.1531. The SLA payload in *IP Endpoint mode* is a proprietary ALBEDO Telecom format. • *All zeroes*: Sets the transmitted pattern to all zeroes.
BERT Patterns	Sets the transmitted and expected test pattern (port A) or the expected test pattern (port B). Supported patterns are: • *PRBS 2^{11}-1 /2^{11}-1 inverted*: This is a pseudo-random bit pattern specified in ITU-T O.153 for error performance measurements below the primary rate (2048 kb/s). The 2^{11}-1 inverted is a 2^{11}-1 bit wise inverted pattern. • *PRBS 2^{15}-1 /2^{15}-1 inverted*: This is a pseudo-random bit pattern specified in ITU-T O.151 for measurements at the primary rate or above. The 2^{15}-1 inverted is a 2^{15}-1 bit wise inverted pattern. • *PRBS 2^{20}-1 /2^{20}-1 inverted*: This is a pseudo-random bit pattern specified in ITU-T O.151 for error performance measurements at the primary bit rate or above. The 2^{20}-1 inverted is a 2^{20}-1 bit wise inverted pattern. • *PRBS 2^{23}-1 /2^{23}-1 inverted*: This is a pseudo-random bit pattern specified in ITU-T O.151 for error performance measurements at the primary bit rate or above. The 2^{23}-1 inverted is a 2^{23}-1 bit wise inverted pattern. • *User*: Sets a 32-bit, user configurable word as the transmitted pattern.

Table 4.3: Ethernet Frame Settings

Setting	Description
User payload	Here it is configured the value of the user payload that is used as the transmitted pattern when *BERT Pattern* is set to *User*.

Some test payloads are byte patterns (*BERT* pattern, a*ll-Zeroes* pattern) but some others have a more complex structure like the *SLA* test payload. Specifically, the SLA test payload used by Ether.Genius / Ether.Sync / Ether.Giga is a proprietary extension of the Operations, Administration and Maintenance (OAM) payload defined by standard ITU-T Y.1731.

4.2. Generation of IPv4 Traffic

Without a Network layer, all the Ethernet traffic generated by Ether.Genius / Ether.Sync / Ether.Giga would be unable lo leave the local network and reach remote networks. The Network Layer, or Layer 3, provides end-to-end connectivity between stations that can use heterogeneous underlying technologies and they are not necessarily attached to the same network. Routers are devices that are designed to manage Layer 3 protocols and data forwarding based on routing tables.

The *Internet Protocol* (IP) is the most popular Layer 3 protocol. It was conceived by the U.S. *Department of Defence* (DoD) during the cold war to facilitate communication between dissimilar computer systems and is a reliable technology. IP interconnects public or private autonomous systems providing a connectionless service.

There are two IP protocol versions (IPv4 and IPv6). IPv4 addresses can be defined as a subset of the IPv6 addressing space but IPv4 and IPv6 can be regarded as different and incompatible network protocols in any other sense. Currently, Ether.Genius / Ether.Sync / Ether.Giga traffic generation functionality is compatible with version four of the IP protocol (IPv4). Traffic analysis include both versions of the IP protocol, IPv4 and IPv6.

The correct operation mode for IPv4 packet generation is the *IP Endpoint* mode. Basically, the traffic generator in *IP Endpoint* mode is configured in the same way than in *Ethernet Endpoint* mode. However, there are some differences to be taken into account:

- In *IP Endpoint* mode, the test equipment becomes a host in an IP network and it has similar properties than any other network equipment. For this reason it is necessary to assign a valid IP profile to the tester either automatically (DHCP) or by hand.
- The test equipment is now ready to use some helper protocols to make the configuration process easier. Specifically, the *Address Resolution Protocol* (ARP), configures destination MAC address without user intervention. The *Domain Name*

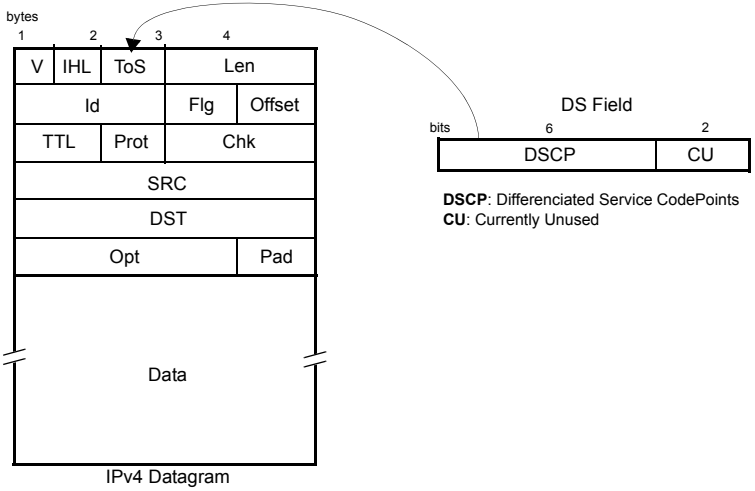

Figure 4.6: IPv4 datagram structure.

Service (DNS) replaces the configuration of the destination IP addresses by the much simpler domain name configuration.

- Ethernet frames carry IPv4 packets with an specific structure and content. It is necessary to configure the IPv4 packet before it is prepared to generate IP datagrams.
- Optionally IPv4 packets carry one or more MPLS labels. The transmission parameters of MPLS labels have to be configured when they are enabled.

4.2.1. Configuring the Physical and MAC Layers

Physical (layer 1) and MAC (layer 2) configuration is similar in *IP Endpoint* and *Ethernet Endpoint* modes (See section 4.1.1, See section 4.1.2). The only difference is that users now have at their disposal the ARP mechanism to configure the destination MAC address automatically. ARP gets the destination MAC address from the network using the destination IPv4 address by means a broadcast protocol.

To use ARP to set the destination MAC address without user intervention, you have to configure the *Destination MAC address from* to ARP (See section 4.1.2). Once ARP has been configured the test unit generates one or several broadcast ARP requests to compute the destination MAC address. Generation of ARP control traffic is automatic and it is not controlled with the RUN button like it happens with the test traffic.

4.2.2. Configuring MPLS

Multi-Protocol Label Switching (MPLS) is a technology designed to speed up IP packet switching in routers by separating the functions of route selection and packet forwarding into two planes:

- *Control Plane*: This plane manages route learning and selection with the help of traditional routing protocols such as *Open Shortest Path First* (OSPF) or *Intermediate System - Intermediate System* (IS-IS).
- *Forwarding Plane*: This plane switches IP packets, taking as a basis short labels prepended to them. To do this, the forwarding plane needs to maintain a switching table that associates each incoming labelled packet with an output port and a new label.

The traditional IP routers switch packets according to their routing table. This mechanism involves complex operations that slow down switching. Specifically, traditional routers must find the longest network address prefix in the routing table that matches the destination of every IP datagram entering the router.

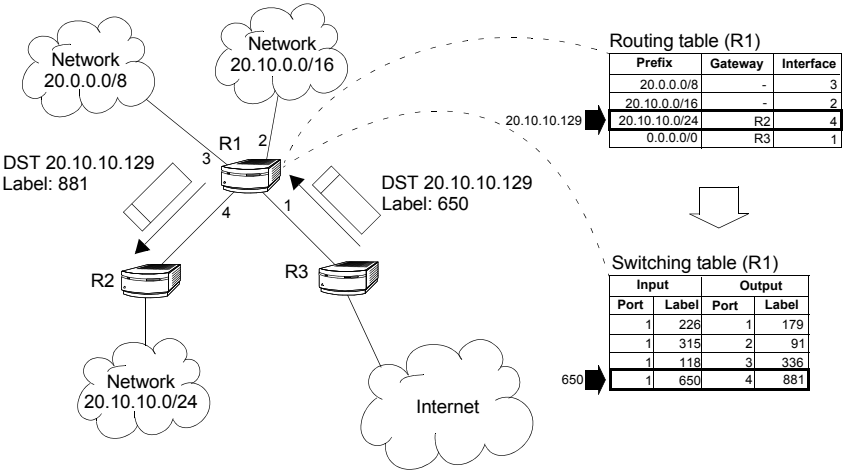

Figure 4.3 Traditional routers have to perform complex operations to resolve the output interface of incoming packets. LSRs resolve the output interface with the help of a simple switching table.

On the other hand, MPLS routers, also known as *Label-Switched Routers* (LSR), use simple, fixed-length label forwarding instead of a variable-length IP network prefix for fast forwarding of packetized data (see Figure 4.3).

MPLS enables the establishment of a special type of virtual circuits called *Label-Switched Paths* (LSP) in IP networks. Thanks to this feature, it is possible to implement resource management mechanisms for providing hard QoS on a per-LSP basis, or to deploy advanced traffic engineering tools that provide the operator with tight control over the path that follows every packet within the network. Both QoS provision and advanced traffic engineering are difficult, if not impossible to solve in traditional IP networks.

To sum up, the separation of two planes allows MPLS to combine the best of two worlds: the flexibility of the IP network to manage big and dynamic topologies automatically, and the efficiency of connection-oriented networks by using preestablished paths to route the traffic in order to reduce packet process on each node.

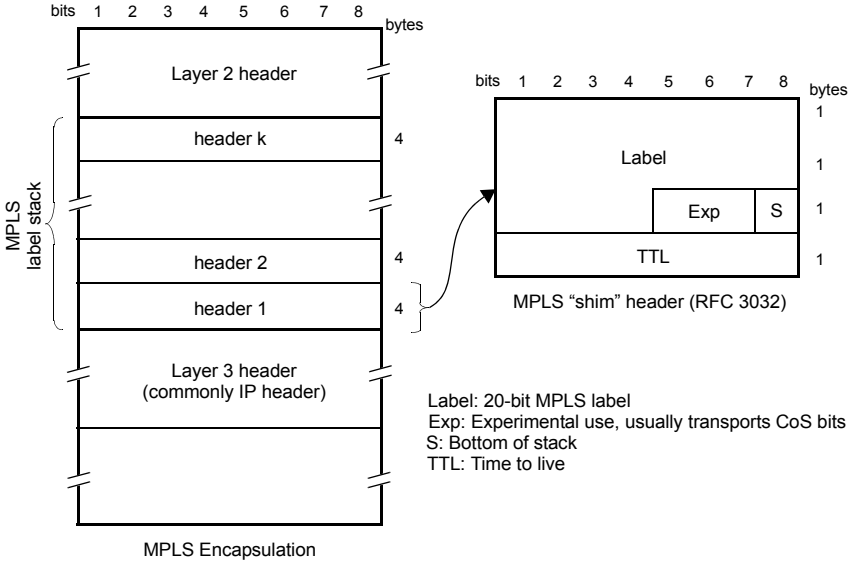

Figure 4.4 MPLS “shim” header format. The label is usually inserted between layer-2 and layer-3 headers.

Ether.Genius / Ether.Sync / Ether.Giga can be configured to generate and analyse MPLS packets carrying one or two labels. The configuration procedure is as follows:

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
3. Select either Port A or Port B to enter in the port specific configuration.
   Note: MPLS configuration settings are not available from Port B because traffic generation is not available from this port.
4. Enter in the *MPLS* menu.
   All settings related with MPLS packet generation configuration are displayed.
5. Select one of the traffic flows between *Flow #1* and *Flow #8* to enter in the flow specific configuration.

6. Configure *Stack configuration* to *Off* to disable MPLS generation over the current flow. Set *Single MPLS label* or *Double MPLS label* to generate IP packets carrying one or two MPLS labels.
7. If you have configured *Single MPLS label* in the previous step, type the correct values for *Bottom label, Bottom traffic* class and *Bottom TTL*. If you have configured *Double MPLS label,* enter the values for *Bottom label*, *Bottom traffic class*, *Bottom TTL*, *Top label*, *Top traffic class* and *Top TTL*.

Table 4.4: MPLS Settings

Setting	Description
Stack configuration	Configures the size of the MPLS label stack for the current flow. Both single and double stack configurations are supported. • *Off*: Disables MPLS generation in the current flow. • *Single MPLS label*: Enables single MPLS label generation in the current flow. The user is expected to configure the bottom label, traffic class and TTL for the MPLS header. • *Double MPLS label*: Enables double MPLS label generation in the current flow. The user is expected to configure both the top and bottom label, traffic class and TTL for the MPLS headers.
Bottom label	MPLS label used for switching traffic in the bottom MPLS header. This field is enabled when the user has selected *Single MPLS label* or *Double MPLS label* in *Stack configuration*.
Bottom traffic class	Contains the traffic class identifier for the bottom MPLS header. It was first thought that this field could carry the 3 Type-of-Service (ToS) bits defined for Class of Service (CoS) definition in the IP version 4, but currently, the ToS field is being replaced by 6-bit *Differentiated Services Code Points* (DSCP). This means that only a partial mapping of all the possible DSCPs into this field is possible. This field is enabled when the user has selected *Single MPLS label* or *Double MPLS label* in *Stack configuration*.
Bottom TTL	This field contains a *Time To Live* value for the bottom MPLS header. The value of this field is decremented by one unit every time the packet traverses an LSR. The packet is discarded if the value reaches 0. This field is enabled when the user has selected *Single MPLS label* or *Double MPLS label* in *Stack configuration*.

Table 4.4: MPLS Settings

Setting	Description
Top label	This field contains the MPLS label used for switching traffic in the top MPLS header. This field is enabled only if the user has selected *Double MPLS label* in *Stack configuration*.
Top traffic class	This field contains the traffic class identifier for the top MPLS header. It was first thought that this field could carry the 3 Type-of-Service (ToS) bits defined for Class of Service (CoS) definition in the IP version 4, but currently, the ToS field is being replaced by 6-bit *Differentiated Services Code Points* (DSCP). This means that only a partial mapping of all the possible DSCPs into this field is possible. This field is enabled only if the user has selected *Double MPLS label* in *Stack configuration*.
Top TTL	This field contains a *Time To Live* value for the top MPLS header. The value of this field is decremented by one unit every time the packet traverses an LSR. The packet is discarded if the value reaches 0. This field is enabled only if the user has selected *Double MPLS label* in *Stack configuration*.

4.4.3. Configuring the Port Local Network Profile

The test equipment requires a local IP profile when it is operating in *IP Endpoint* mode. Even if the traffic generator has been configured to work with a custom IP address (different to the local IP address), the equipment still requires an internal address for some tests like the IP Ping or the Traceroute. Furthermore, some control and signalling protocols may work with the information stored in the local IP profile.

Configuration of the local IP profile is available form the port specific settings within the Setup menu (See section 2.3).

4.4.4. Configuring the IPv4 Datagram

The IPv4 packet content is set much in the same way that the MAC frame content. However, in this case, MAC addresses are replaced by IPv4 addresses. Of course,

IPv4 datagrams have their own structure and they contain some fields not present in Ethernet frames.

Table 4.5: IPv4 Packet Settings

Setting	Description
Source IPv4 address from	Establishes the origin of the source IPv4 address for the current stream. There are two possible settings: • *Local*: The source address is set to the IPv4 address configured in the port local profile. The local address may be either configured by means the DHCP protocol or in may be static. • *Manual*: The source address is set to the value configured in *Source IPv4 address*. Use manual IPv4 addresses if you want to simulate traffic generated by an equipment different to the tester or, in multi-stream operation, to simulate traffic transmitted from different hosts. Probably, you will want to avoid duplicated IP addresses in your network. For this reason, make sure that no other equipment is using the manually configured IPv4 address.
Source IPv4 address	Source IPv4 address carried by the packets generated in the current stream if *Source IPv4 address from* is set to *Manual*. The address is entered in decimal, four-dotted format. Any address between 0.0.0.0 and 255.255.255.255 is admitted as a source IPv4 address.
Destination IPv4 address from	Establishes the origin of the destination IPv4 address for the current stream. There are three different settings available for configuration: • *Manual*: The destination address is set to the value configured in *Destination IPv4 address*. • *Range*: Test data in the current stream is transmitted to a group of IPv4 addresses configured with *Destination IPv4 address* and *Address range size*. Use this option if you want to deliver the test data sequentially to many different destinations.
Destination IPv4 address from	• *Host name*: Uses the Domain Name Service (DNS) to set the destination IP address by using descriptive alphanumeric strings. The DNS mechanism requires intervention of at least one DNS server. The DNS server IP address has to be configured in the local port profile either statically or by means DHCP.

Table 4.5: IPv4 Packet Settings

Setting	Description
Destination IPv4 address	Destination IPv4 address carried by the packets generated in the current stream if *Destination IPv4 address from* is set to *Manual*. The address is entered in decimal, four-dotted format. Any address between 0.0.0.0 and 255.255.255.255 is admitted as a destination IPv4 address.
Destination IPv4 address (DNS)	Destination IPv4 address carried by the packets generated in the current stream if *Destination IPv4 address from* is set to *Host name*. This is a read only field that it cannot be edited directly. It displays the result of the DNS name resolution carried out with the host name configured in *Destination host name.*
Address range size	Configures the number of IPv4 addresses within an address range. This control is valid only if *Destination type* is set to *Range*. In this case, the IP datagrams transmitted in the current stream will contain as many destination addresses as previously configured in this field. The destination IP address is increased by one unit for each transmitted frame starting with the value configured in *Destination IPv4 address*. If there are no more addresses left in the range, transmission returns to the initial address and starts the process from the beginning. Transmission of destination IPv4 address ranges is compatible with transmission of destination MAC address ranges but the address number of the MAC address range is always fixed to the same number that the IP range. It is not possible to transmit a destination MAC address range with a single IPv4 address.
Destination host name	Domain name to be used as a destination if *Destination type* is set to *Domain name*. Unlike IP addresses, domain names are easy-to-remember alphanumeric strings but they have to be translated to IP addresses before any packet can be sent to the destination. The translation process requires the intervention of at least one DNS server. The DNS server IP address has to be configured in the local port profile either statically or by means DHCP.

Table 4.5: IPv4 Packet Settings

Setting	Description
DSCP	Differentiated Services Code Point. It is 6-bit class of service (CoS) field defined to set packet groups with different priorities or to provide specific treatments to special packets within a network or an administrative domain. Any value from 0 to 63 is allowed for this field. Specific actions to be carried out on frames with different DSCPs depend on the network and the service provider.
TTL	Initial Time To Live value configured in the packets transmitted in the current stream. The TTL is decreased by one unit each time it lefts a network node. If the value reaches zero, then the packet is discarded. The TTL is then a measure of the number of nodes the packet is allowed to transverse before reaching its destination.
UDP	Enables or disables transmission of the User Datagram Protocol (UDP) in the current stream. The UDP is defined in RFC 768 and it is a lightweight transport protocol for unreliable data transmission. RFC 768 defines an eight-byte fixed length header for UDP that is generated when UDP generation is enabled in the stream. If UDP generation is on, the *Transport protocol* field is set to 17. This value cannot be edited by the user.
Transport protocol	This setting contains an 8-bit word that constitutes the protocol identifier to be transmitted by the traffic generator. TCP uses 6 as the protocol number, UDP uses 17 for the same purpose and ICMP uses number 1. However, the payload structure does not match the structure corresponding to these protocols even if the correct protocol number is configured. To enable UDP header and payload generation, enable UDP in the current stream.
Source port	Source transport layer port transmitted in the UDP header in the current stream. Ports are service identifiers used to multiplex data from different applications generated by IP hosts. The tester supports source port generation for UDP streams only.

Table 4.5: IPv4 Packet Settings

Setting	Description
Destination port	Destination transport layer port transmitted in the UDP header in the current stream. Ports are service identifiers used to multiplex data from different applications generated by IP hosts. The tester supports destination port generation for UDP streams only.

All the IPv4 datagram configuration lays within the *Network* menu. The network menu is not enabled unless the port is configured in TX / RX mode. For this reason, there is no Network configuration menu in Port B (Port B does not support TX / RX mode). The procedure to follow to configure the IP datagram is described below:

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1). Check that your tester is operating in *IP Endpoint* mode (See section 2.1) and that the port is in *TX / RX*.
2. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
3. Select Port A to enter in the port specific configuration.
   Note: There is no network configuration for Port B.
4. Enter in the *Network layer* menu.
5. Select one of the traffic flows between *Flow 1* and *Flow 8* to enter in the flow specific configuration.
   All settings related with network configuration in the current flow are displayed.
6. Enter the source IP address with the help of the *Source IPv4 address from* and *Source IPv4 address* controls. You can configure the IP address from the local IP profile as the source address or enter a custom address.
7. Enter the destination IPv4 address or addresses by using the *Destination IPv4 address from*, *Destination IPv4 address*, *Address range size* and *Destination host name*. If you choose to generate a destination address range you will be requested to enter the number of addresses that made up the range. If you choose to enter the destination as a host name rather than an IPv4 address, you will be requested to enter a valid domain name.
8. Configure the DSCP and TTL if necessary.
9. Enable or disable UDP generation and analysis with the help of the *UDP* control.
10. If you have enabled UDP, enter the *Source Port* and *Destination Port* to be used in the generated UDP packets.
11. If you have not enabled UDP, configure the *Transport Protocol* code.
12. If necessary, repeat the IPv4 configuration process for one or more traffic flows (*Flow #1* to *Flow #8*) available from the *Network layer* menu.

4.4.5. Setting the Bandwidth Profile

Setting the bandwith profile in IP Endpoint mode is the same that in Ethernet Endpoint mode (See section 4.1.3)

4.4.6. Choosing the Test Payload for IPv4

Ether.Genius / Ether.Sync / Ether.Giga, include special packet payloads and patterns required for all usual applications, including BER tests and SLA tests. Payloads and patterns available in *IP Endpoint* mode are similar than in *Ethernet Endpoint* mode (See section 4.1.4). However, there is a difference concerning the SLA payload. While in Ethernet Endpoint the SLA payload is defined as an extension of the ITU-T Y.1731 structure, in *IP Endpoint*, this payload is ALBEDO Telecom proprietary. In practical terms, the new structure of the SLA payload should not make any difference.

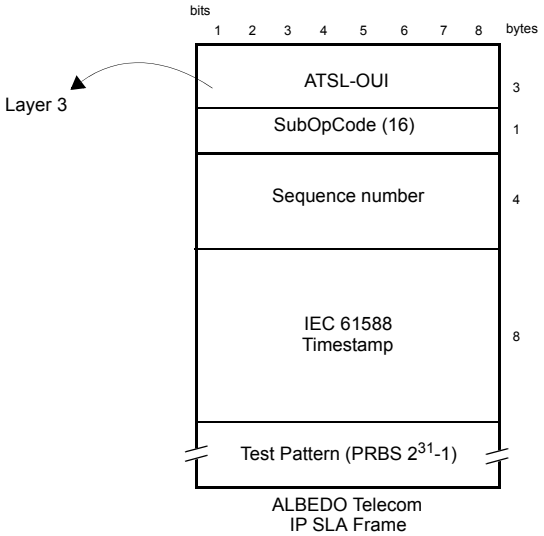

Figure 4.7: Albedo payload for SLA tests (IP Endpoint mode).

4.5. Event Insertion

Sometimes it is necessary to insert events in the generated signal to stress the DUT/ SUT. Ether.Genius / Ether.Sync / Ether.Giga implements extended event insertion capabilities. The procedure to set event insertion with the your test unit is as follows:

1. From the *Home* panel, go to *Test*,
   The *Test* configuration panel is displayed.

2. Select *Insertion*
   The event insertion menu is displayed.
3. Select the event to be inserted with the help of the *Event to be inserted* menu item.
4. Select the insertion mode for the event selected in the previous step with the help of the *Mode* menu item. Available insertion modes are: *Single*, *Rate*, *Burst* or *Random*.
5. Configure the insertion parameters with the help of the *Event rate*, *Number of frames*, *Probability (%)* and *Frame size (bytes)* menu items.
6. Start insertion by pressing the EVENT button.
   Note: Depending on the insertion mode, event insertion will finish automatically or you will need to press EVENT a second time to stop.

Table 4.6: Event Insertion Settings

Setting	Description
Event to be inserted	Contains a selection list with events the user can choose for insertion in the transmitted data stream.The events available are: • *None*: Disables event insertion in the target port and flow. No event will be generated in case the user presses EVENT. • *FCS*: Generates *Frame Check Sequence* errors. A frame with a FCS error is a frame with a legal size which contains an invalid FCS field. In normal circumstances, FCS errors are caused by transmission errors. An optical Ethernet link with a poor power budget may experience FCS errors • *IPv4 checksum*: Generates frames with an invalid IPv4 header checksum. In normal circumstances, in Ethernet networks, IP checksum errors are related with corrupted traffic generation. IPv4 checksum error insertion is available only in *IP endpoint* mode. • *Undersized frames*: This event is used to generate frames shorter than the minimum legal size (64 bytes). The frame size of an undersize frame is configured through the *Frame size (bytes)* field. • *TSE*: Generates *Test Sequence Errors*, One TSE is equivalent to a single bit difference between the transmitted and the received test pattern (PRBS or other). TSE event insertion is not available if the target port transmitter is not configured for transmission of a BER test pattern.

Table 4.6: Event Insertion Settings

Setting	Description
Mode	Configures the way events are inserted in the outgoing signal. Depending on whether the insertion event is an anomaly or a defect there are different insertion modes. For anomalies, the insertion modes are: • *Single*: A single event is inserted. Event insertion is triggered when the EVENT key is pressed. • *Burst*: A burst events a configurable number of events is inserted. Burst start is triggered with the EVENT key. • *Rate*: Events are inserted with a configurable rate. Insertion is deterministic (the time interval between consecutive events is a constant). Insertion starts if the EVENT key is pressed. Insertions stops when the EVENT key is pressed again. • *Random*: The number of events generated is random.For each insertion opportunity, the transmitter decides if the event is inserted or not depending on a user configurable insertion probability.
Event rate	If the insertion mode has been set to *Rate*, this fields sets the rate at which events are inserted in the outgoing signal. The rate is entered in scientific notation: A x 10^{-B} In this notation B is a number between -3 and -9 (both included) and A is a real positive number smaller than 10.
Number of frames	If insertion mode has been set to *Burst*, this field sets the number of events that makes up the burst. For example a burst of 10 bit errors is made of ten consecutive TSE errors.
Probability (%)	If the insertion mode is set to *Random,* this is the probability of a single event occurrence expressed as a percentage.
Frame size (bytes)	If *Event to be inserted* has been configured to *Undersized frames*, this field configures the frame length corresponding to these undersized frames expressed in bytes. It is a number between 32 and 63.
Target flow	Sets the traffic flow where the event is going to be inserted. Currently it is only supported insertion in flow 1.

Chapter 5
Basic Frame Analysis

The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga can be used to get advanced traffic counts and statistics about Ethernet and IP networks operating at rates up to 1 Gb/s. These statistics include frame and error counts, bandwidth statistics, quality of service statistics, frame size statistics and other results.

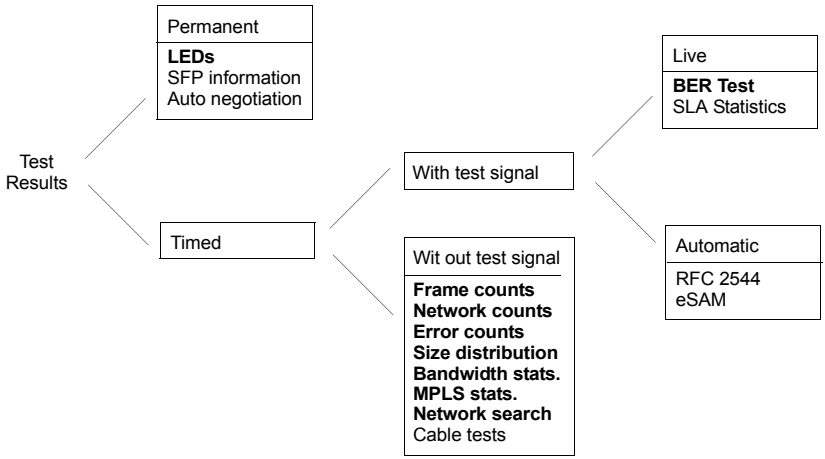

Figure 5.1: Statistics and counts available in Ether.Genius / Ether.Sync /Ether.Giga.

Test results supplied by Ether.Genius / Ether.Sync / Ether.Giga can be classified in many different ways. One of them is to think in test results that require a previous test start action with the RUN button (timed results) or results which are always available when the equipment is on (permanent results). A second approach is to classify test results in those which depend on an special test signal transmitted from an special traffic generator and results available in any case, even if there is no test signal available for analysis.

This chapter deals mainly with timed results which does not require test signals and the closely related LED results even if these are permanent results. At the end of the chapter BER test results are also addressed, but these are different in nature to all other measurements because they require a BERT payload / pattern received in the test interface. These test may be generated by the same Ether.Genius / Ether.Sync / Ether.Giga unit or a remote tester.

5.1. Global Counts and Statistics

Global frame counts and statistics are those not associated to any particular stream. Some global statistics have a per-stream statistic counterpart. Examples of this are the bandwidth statistics and some frame counts.

Global counts and statistics for Port A and Port B are identical. Global frame statistics are controlled by the RUN button. That means that results are not collected if a test is not started before. Once the test is running results are upgraded in real time.

Counts and statistics described in this section are available both in endpoint (*Ethernet Endpoint*, *IP Endpoint*) and through mode.

5.1.1. Frame Counts

Frame counts provide information about how many frames have been received in the test interface from the beginning of the test. These counts are also useful to classify the frames received in different families.

To display the transmitter statistics follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *Frame layer statistics*.
4. Check the *TX frames*, *RX frames*, *Unicast frames*, *Multicast frames*, *Broadcast frames*, *VLAN*, *IEEE 802.1ad*, *Q-in-Q*, *Control frames*, *Pause frames* and *BPDU* counters.

Table 5.1: Global Frame Statistics

Metric	Description
TX frames	Total number of frames transmitted by one tester port since the test started.
RX frames	Total number of frames received by one tester port since the test started.

Table 5.1: Global Frame Statistics

Metric	Description
Unicast frames	Total number of Ethernet unicast frames received from the beginning of the test. Unicast frames are recognised because they contain a unicast destination MAC address. Unicast MAC frames have their multicast bit set to '0'. The multicast bit of a MAC address is the least significant bit of the more significant address byte.
Multicast frames	Received Ethernet multicast frames from the beginning of the test. Ethernet multicast frames have their multicast bit in their destination MAC address set to '1'. The multicast bit of a MAC address is the least significant bit of the more significant address byte.
Broadcast frames	Total number of Ethernet broadcast frames received from the beginning of the test. Broadcast frames carry the broadcast Ethernet address (*FF:FF:FF:FF:FF:FF*) in the destination field.
VLAN	Total number of Ethernet VLAN frames transmitted from the beginning of the test. IEEE 802.1Q VLAN frames contain an special Ethertype (Type / Length field) value (*0x8100*).
IEEE 802.1ad	Total number of *Provider Bridge* (PB) frames received from the beginning of the test. PB frames are defined by standard IEEE 802.1ad. PB frames contain two VLAN tags referred as C-VLAN and S-VLAN. The C-VLAN carries Type *0x8100*. The S-VLAN has the special Type *0x88a8*.
Q-in-Q	Total number of double-tagged VLAN frames received from the beginning of the test. Q-in-Q frames contain an S-VLAN and a C-VLAN but they are not compliant with the IEEE 802.1ad standard. This standard requires the Type field for the S-VLAN to be *0x88a8* but not-standard Q-in-Q frames use different values like *0x8100*, *0x9100*, *0x9200* or *0x9300*.

Table 5.1: Global Frame Statistics

Metric	Description
Control frames	Total number of Ethernet MAC control and supervision frames received from the beginning of the test. Ethernet control frames are recognised due to an special Ethertype (Type / Length field) value (*0x8808*).
Pause frames	Total number of Ethernet *Pause* frames received from the beginning of the test. Pause frames are an special type of control frames and therefore their Ethertype is 0x8808. The specific features of *Pause* frames is that their *Opcode* field is 0x0001 and their destination MAC address is *01:80:C2:00:00:01* (a multicast MAC address).

5.1.2. Error Counts

Ether.Genius / Ether.Sync / Ether.Giga are prepared to get any defect or fault in the received data stream. These faults include invalid checksum, alignment or size and frame structure defects. If configured in pass-through mode, Ether.Giga / Ether.Sync / Ether.Genius, automatically drops frames with errors but these are still counted and presented in the error statistics panels. To access to the error statistics:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *Error statistics*.
4. Check the *FCS*, *Undersized*, *Oversized*, *Jabbers*, *IPv4 / IPv6 errors* and *UDP errors* counters.

Table 5.2: Global Error Counts

Metric	Description
FCS	Count of all the FCS errors detected from the beginning of the test. A frame with a FCS error is a frame with a legal size which contains an invalid FCS field. FCS errors are caused by transmission errors. An optical Ethernet link with a poor power budget may experience FCS errors
Undersized	Total number of received frames which are smaller than 64 bytes.
Oversized	Total number of received frames which are larger than the configured MTU.

Table 5.2: Global Error Counts

Metric	Description
Jabbers	Jabber count from the beginning of the test. A Jabber is defined as a frame greater than 1518 bytes with a bad CRC.
IPv4 / IPv6 errors	Displays information about the total number of errored IPv4 and IPv6 packets. An IPv4 or IPv6 packet is declared an errored packet if it is improperly built: Datagram header has an unexpected structure, field values are different to the expected ones or it has an unexpected length. An IPv4 datagram is also considered to be invalid if it contains checksum errors.
UDP errors	Displays information about the total number of errored UDP packets. UDP packets are considered to contain errors if either their structure or their content is invalid. For UDP, this means that the packet has an unexpected length or it contains checksum errors.

5.1.3. Network Counts

The Network layer statistics panel has similar purpose than the Frame layer statistics but in this case it displays counters related with the IP layer rather than with Ethernet. Network statistics are not available in *Ethernet Endpoint* mode because the IP structure is neither generated nor decoded in this operation mode.

Table 5.3: Global Network Statistics

Metric	Description
IPv4 / IPv6 TX	Aggregated number of transmitted IPv4 and IPv6 packets since the last test start. IPv4 packets are encapsulated in Ethernet frames (with or without an LLC / SNAP header) carrying the 0x0800 Type. IPv6 packets are encapsulated (with or without the LLC / SNAP header) with the Ethertype field set to 0x86DD.
IPv4 RX	Total number of received IPv4 packets since the last test start. IPv4 packets are encapsulated in Ethernet frames (with or without an LLC / SNAP header) carrying the 0x0800 Type.

Table 5.3: Global Network Statistics

Metric	Description
IPv6 RX	Total number of received IPv6 packets since the the beginning of the test with RUN. IPv6 packets are encapsulated in Ethernet frames (with or without an LLC / SNAP header) carrying the 0x86DD Type.
Unicast pkts.	Aggregated count of unicast IPv4 and IPv6 packets received from the beginning of the test. An IPv4 unicast packet is a packet directed to a unicast IPv4 address. An IPv4 unicast address is any valid, not-broadcast Class A (1.0.0.1 - 126.255.255.254), Class B (128.1.0.1 - 191.255.255.254) or Class C (192.0.0.1 - 192.255.254.254) address. An IPv6 unicast packet is a packet directed to a unicast IPv6 address. All non-multicast IPv6 (prefix ff00::/8) packets are considered to be unicast.
Multicast pkts.	Aggregated count of unicast IPv4 and IPv6 packets received from the beginning of the test. An IPv4 multicast packet is a packet directed to a multicast IPv4 address. An IPv4 multicast address is any valid Class D (224.0.0.0 - 239.255.255.255) address. An IPv6 multicast packet is a packet directed to a multicast IPv6 address. IPv6 multicast addresses is an address starting with the ff00::/8 prefix.
Broadcast pkts.	Total count of broadcast IP packets received from the beginning of the test. An IPv4 broadcast packet is a packet directed to the currently configured network broadcast address or the global broadcast address (255.255.255.255). The network broadcast address is the IPv4 address that has all the host bits set to 1. There is no definition for IPv6 multicast addresses. For this reason IPv6 statistics are not included in this result field.
UDP packets	UDP packets or segments are IP packets carrying the User Datagram Protocol (UDP) defined in RFC 768. The protocol number assigned to UDP is 17.

Table 5.3: Global Network Statistics

Metric	Description
ICMP / ICMPv6	Aggregated count of ICMP and ICMPv6 messages received from the beginning of the test. ICMP packets carry Internet Control Message Protocol (ICMP) messages defined in RFC 792. ICMPv6 messages carry the version 6 of Internet Control Message Protocol defined in RFC 4443. The protocol number assigned to ICMP is 1. For ICMPv6 the protocol number is 58.

To display the network statistics corresponding to the last (or current) test, follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *Network layer statistics*.
4. Check the *IPv4 TX*, *IPv4 RX*, *Unicast pkts.*, *Multicast pkts.*, *Broadcast pkts.*, *UDP* and *ICMP* counters.

5.1.4. Bandwidth Statistics

Bandwidth statistics inform about how many frames and bits are you receiving per time unit and how much of the available bandwidth is being used by the traffic. Traffic statistics are supplied at many different transmission layers, including Ethernet, IP and UDP. IP and UDP is closer to the amount of usable data and Ethernet statistics are more related with the bandwidth available for transmission. For this reason, the Ethernet bandwidth is expressed as a percentage of the nominal transmission bandwidth but this is not the case with IP and UDP bandwidth statistics.

Bandwidth statistics are timed measurements, you need to run a test to start collecting results. In order to access to these results follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *Bandwidth statistics*.
4. Select *General statistics*.

5. Check the *Eth. (current)*, *Eth. (min.)*, *Eth (max.)*, *Eth. unicast*, *Eth. multicast*, *Eth. broadcast*, *IPv4 (current)*, *UDP (current)* results.

Table 5.4: Bandwidth Statistics

Metric	Description
Eth. (current)	Total amount of Ethernet traffic received during the last second computed in bits per second, frames per second and as a percentage of the total channel capacity. The frame bits considered in Ethernet traffic statistics are between the first bit of the destination address field to the last bit of the FCS field. The current Ethernet traffic counter requires previous initialization of a test with the RUN button.
Eth. (max)	Peak value of Ethernet traffic registered from the beginning of the test. The *Ethernet (max)* value displays the maximum value found in the *Ethernet (current)* field since the test started. The *Ethernet (max)* is displayed in three different units: bits per second, frames per second and percentage of the overall channel capacity.
Eth. (min)	Minimum value of Ethernet traffic registered from the beginning of the test. The *Ethernet (min)* value displays the minimum value found in the *Ethernet (current)* field since the test started. The *Ethernet (max)* is displayed in three different units: bits per second, frames per second and percentage of the overall channel capacity.
Eth. unicast	Total amount of unicast Ethernet traffic received during the last second computed in frames per second. Unicast frames are recognised because they contain a unicast destination MAC address. Unicast MAC frames have their multicast bit set to ‘0’. The multicast bit of a MAC address is the least significant bit of the more significant address byte.
Eth. multicast	Total amount of multicast Ethernet traffic received during the last second computed in frames per second. Ethernet multicast frames have their multicast bit in their destination MAC address set to ‘1’. The multicast bit of a MAC address is the least significant bit of the more significant address byte.

Table 5.4: Bandwidth Statistics

Metric	Description
Eth. broadcast	Total amount of broadcast Ethernet traffic received during the last second computed in frames per second. Broadcast Ethernet frames carry the broadcast Ethernet address (*FF:FF:FF:FF:FF:FF*) in the destination field.
IPv4 (current)	Total amount of IPv4 traffic received during the last second computed in bits per second, frames per second and as a percentage of the overall channel capacity. The current IPv4 traffic counter requires previous initialization of a test with the RUN button.
IPv6 (current)	Total amount of IPv6 traffic received during the last second computed in bits per second, frames per second and as a percentage of the overall channel capacity. The current IPv6 traffic counter requires previous initialization of a test with the RUN button.
UDP (current)	Total amount of traffic associated to the Ethernet / IP / UDP payloads computed in bits per second, frames per second and as a percentage of the overall channel capacity. The current *User traffic* counter requires previous initialization of a test with the RUN button.

5.1.5. Frame Size Distribution

Frame size is important because it tells how a network is used. Some applications, like VoIP use short frames while most data applications based on a client / server architecture use short frame lengths for the client requests and long frames for the server replies. The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga provide frame size results as described in standard RFC 2819. The procedure for displaying the received frame size distribution is as follows:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *Frame size histogram*.
4. Check the frame size intervals: *64 or less*, *65 - 127*, *128-255*, *256 - 511*, *512-1023*, *1024 - 1518, 1519 - 1522*, *1523-1526*, *1527 - MTU*.

5.1.6. MPLS Statistics

Network statistics and all other results are not affected when the received traffic contains MPLS labels. For example, IPv4 datagrams are still IPv4 packets when they have one or more MPLS labels and they are still recognised as unicast, multicast or

broadcast packets when this happens.MPLS test results are limited to an indication of the presence of MPLS in the received traffic and some statistics about the MPLS stack size. To display the MPLS statistics follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *MPLS statistics*.
4. Check the *MPLS labels per packet (min.)* and *MPLS labels per packet (max.)* counters.

Table 5.5: MPLS statistics

Metric	Description
MPLS labels per packet (min.)	Minimum number of MPLS labels per packet found in the received traffic from the beginning of the test. If it has been received at least one frame without MPLS labels, this result will be set to 0.
MPLS labels per packet (max.)	Maximum number of MPLS labels per packet found in the received traffic from the beginning of the test.

5.2. Using the Network Search Capability

Network Search is an optional monitoring tool for Ether.Genius / Ether.Sync / Ether.Giga that reports the most popular MAC, IPv4 and IPv6 addresses found in the network. Also, if you are connected to a tagged interface, the network search can be configured to collect the most viewed VLANs. Network search can be used to look for an specific traffic flow in your network or maybe as a preliminary analysis tool before filtering the interesting traffic flows and getting detailed statistics about them. In fact, the Network Search capability is prepared to be used as a preliminary analysis tool: Once the traffic search has finished the user can choose which filters to configure following the results of the network search. All this is done by a simple key press. Network Search is used in the following way:

1. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
2. Select either *Port A* or *Port B* to enter in the port specific configuration.
3. Go to *Search network by* and define your search field: *Source MAC address*, *Destination MAC address*, *VLAN, Source IPv4 address*, *Destination IPv4 address, Source IPv6 address, Destination IPv6 address*.
   Note: *Source IPv4 / IPv6 addresses* and *Destination IPv4 / IPv6 addresses* are available only if you are working in an IP mode (*IP endpoint* or *IP through*). You can configure *VLAN* even if you are not connected to a tagged interface but you will not collect any VLAN information when you start the measurement. In order to

disable the network search in the current port set *Search network* by to *Off*.
Note: You can run one Network search in Port A and one in Port B. To do that, simply configure a search field for each port.

4. Leave the previous panel to *Home* and go to *Results*.
   The test port results panel is displayed.
5. Select either *Port A or Port B* to enter in the port specific results.
6. Go to *Network Search*.
7. Start a test by pressing RUN.
   The network search panel is filled with popular addresses or VLANs found in your network. There is a percentage that indicates how many times each address or VLAN has been found to the time.
8. Finish the search by pressing RUN a second time.
9. Use the cursors and the ENTER button to choose which addresses or VLANs will be used to configure the filters on the receiver.
10. Press the *Configure* contextual button to configure the filters with the previously collected data. To get extended information about these filters you will need to run a second measurement once they have been configured with this action (See chapter 6).

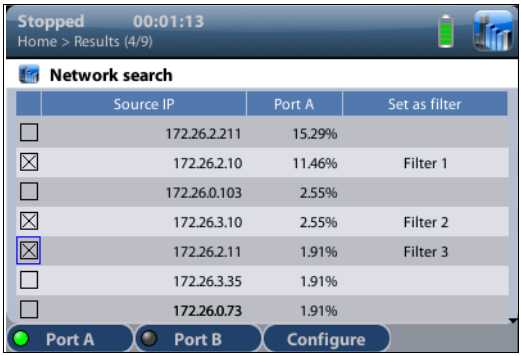

Figure 5.2: Ether.Genius / Ether.Sync / Ether.Giga network search capability

5.3. The LEDs Panel

The LEDs panel offers a quick view of the current Ether.Genius / Ether.Sync / Ether.Giga connection and operation status. They are permanent indicators. That means that no test has to be started to get the information from the LEDs.

There are two hardware global summary LEDs in the equipment (one for Port A and one for Port B), six summary LEDs for each test port (*Link*, *Traffic*, *Frame*, *VLAN,*

Network and *Pattern*). These summary LEDs summarize the information of the events shown in the LEDs panel. To display the LEDs panel use the LEDS key. If the LEDs panel is already visible press LEDS again to return to the previous screen.

The LEDs have two operation modes:

- *Live*: Events are shown in real time. If something happens the corresponding LEDs change their colour to signal the event. LEDs return to their original status once the event disappears.
- *History*: The LEDs keep their original Anomaly / Defect status when the event disappears. This is useful when the tester is left a long time under operation and the user wants to receive quick feedback of past events.

The live or history modes can be configured from the LEDs panel by means the contextual keyboard. The *History* (F3) contextual button sets or unsets the history mode. If the history mode is enabled, then the Reset (F4) button resets the LEDs history.

Figure 5.3: Albedo Ether.Giga LEDs panel.

Orange is used to signal anomalies and red indicates defects but there are more possible LED status:

- : OK, the event or events that correspond with the LED are not found in the incoming signal.
- and : This is the colour displayed if faulty conditions are found in the signal. Conditions marked with tend to be more important than the ones marked with .
- : Shows that no operation condition can be established due to the lack of matching traffic for the corresponding event. For example, the FCS led is if no traffic is received because there are no frames where to check the FCS. It can

also indicate that the LED has been disabled due to the presence of a more important event.

Table 5.6: LED Indications

Metric	Description
REF	External clock input status LED. Displays green colour if a valid reference signal is found in the corresponding clock input port. Otherwise, the displayed LED colour is red. The *REF* LED applies to Ether.Genius and Ether.Sync only.
LOCK	This LED is green when synchronization with the current clock has succeed. If external synchronization fails for some reason, the LED colour changes to red. Sometimes, it may take a few seconds to achieve synchronization with the external clock input. The LED colour during this transient period is yellow. The *LOCK* LED is enabled only when the clock reference is detected in the corresponding reference input clock interface. The *LOCK* LED applies to Ether.Genius and Ether.Sync only.
1000	The port is operating at 1000 Mb/s either from the optical or the electrical port. Port speed is decided immediately after connecting the port to the DUT / SUT using Ethernet auto-negotiation or it is forced by the user.
100	The port is operating at 100 Mb/s from the electrical port. Port speed is decided immediately after connecting the port to the DUT / SUT using Ethernet auto-negotiation or it is forced by the user.
10	The port is operating at 10 Mb/s from the electrical port. Port speed is decided immediately after connecting the port to the DUT / SUT using Ethernet auto-negotiation or it is forced by the user.
RX	At least one frame was received during the current second in the current interface.
TX	At least one frame was transmitted during the current second in the current interface.

Table 5.6: LED Indications

Metric	Description
FCS	At least one frame with FCS errors have been found during the current second. A frame with a FCS error is a frame with a legal size which contains an invalid FCS field. FCS errors are caused by transmission errors. An optical Ethernet link with a poor power budget may experience FCS errors
Jabber	At least one jabber was received during the current second. Jabbers are defined as frames greater than 1518 bytes with a bad CRC.
UnderS	At least one undersized frame was received during the current second. An undersized frame is a frame which has a size smaller than 64 bytes.
OverS	At least one oversized frame was received during the current second. An oversized frame is a frame which has a size larger than the configured MTU.
VLAN	At least one frame with an VLAN tag was received during the current second in the current interface.
802.1ad	At least one frame containing the IEEE 802.1ad Ethertype value (*0x88a8*) has been detected during the current second in the current interface. IEEE 802.1ad frames carry two VLAN tags known as S-VLAN and C-VLAN.
Q-in-Q	At least one frame carrying two VLAN tags but not the IEEE 802.1ad Ethertype (*0x88a8*) has been detected during the current second in the current interface. Q-in-Q frames carry two VLAN tags known as S-VLAN and C-VLAN but they are not compliant with any international standard like the IEEE 802.1ad.

Table 5.6: LED Indications

Metric	Description
IPv4/v6	If this led is green, at least one correct IPv4 or IPv6 datagram was received during the current second. If red, this led indicates that one incorrect datagram has been received during the current second. One IPv4 / IPv6 datagram is incorrect if its header has an unexpected structure, field values are different to the expected ones, it has an unexpected length or it contains checksum errors.
UDP	If this led is green, at least one correct UDP packet was received during the current second. If red, this led indicates that one incorrect datagram has been received during the current second. One UDP packet is incorrect if it has an unexpected length or it contains checksum errors.
LSS	Loss of Sequence Synchronization. This event indicates that the expected test pattern does not match the actually received test pattern. This event does not apply if the tester is configured to receive an SLA payload and, for framed analysis, it is referred to the stream 1 (the only stream with pattern analysis capabilities).
TSE	Test Sequence Error. One TSE is equivalent to a single bit difference between the transmitted and the received test pattern (PRBS or other). This event does not apply if the tester is configured to receive an SLA payload and, for framed analysis, it is referred to the stream 1 (the only stream with pattern analysis capabilities).

5.4. The Event Logger

Global counts, statistics and LEDs provide information about which events and how many of them have been registered but they do not say too much about how they are distributed in time. These information is supplied by the Ether.Genius /Ether.Sync / Ether.Giga graphical representation tool or *event logger*.

With the Help of the event logger function, you can select one or various events and trace them so that all changes along with the time and date these changes are registered are recorded with a 1 second resolution. The event logger provides different representations and different zoom levels to enable event analysis at different time scales.

5.4.1. Configuring the Event Logger

Traceable Events are categorized in different classes. Moreover, each test port has its own traceable events. These events may be different for each test port.

Table 5.7: Logging event categories

Event class	Description
Anomalies / defects	Accounts for Ethernet /IPv4 /IPv6 anomalies and defects. This category includes the following events: *Link*, *FCS*, *Undersized*, *Oversized*, *Jabbers*, *IPv4 / IPv6 errors*, *UDP errors.*
Bandwidth statistics	Reports bit rates in *bit/s* for different protocol layers, including: *Ethernet bit rate*, *IPv4 bit rate*, *IPv6 bit rate*, *UDP bit rate*. The *Bandwidth statistics* category includes one set of traceable events for each filter (*Filter 1*, *Filter 2*,...) and one additional global statistics subclass.
Frame layer statistics	Event category that includes transmission statistics related with different frame structures: *RX frames*, *TX frames*, *Unicast*, *Multicast*, *Broadcast*, *VLAN*, *IEEE 802.1ad*, *Q-in-Q*, *Control*, *Pause*.
SLA statistics	Provides information about, delay, delay variation, frame loss and other QoS parameters. Includes the following events*: RX frames*, *RX bytes*, *FTD*, *FDV*, *Lost frames*, *Duplicated frames*, *SES*. This category has one subclass for each filter (*Filter 1*, *Filter 2,...*) so that metrics from different frame flows can be simultaneously traced and compared.
BERT	Reports BER results by means the *LSS*, *TSE* traceable events.
Synchronization	This event category includes all events related with PTP and Synchronous Ethernet. Both message statistics, and time / frequency metrics are available. The traceable event list is: *Sync*, *Delay request*, *Delay response*, *Peer delay request*, *Peer delay response*, *Follow-up*, *Announce*, *Signaling*, *Management*, *Domain mismatch*, *Sync PTD*, *Sync PDV*, *Delay req. PTD*, *Two-way PTD*, *Sync IAD*, *PTD slave frequency offset*, *PTD slave phase offset*, *Frequency deviation*.

To enable the Event logger follow these steps.

1. From the *Home* panel, go to *Test*,
   The *Test* configuration panel is displayed.

2. Select *Event logger setup*.
   The event logger configuration menu is displayed.
3. Enable event logging with the help of the *Enable* control.
4. Optionally, clear the currently selected filters with the *Clear all filters* menu.
5. Select the *Port A* or *Port B* (if you are working with Ether.Genius you can also select *Port C* depending on the current operation mode)
6. Choose the event categories corresponding to the events you want to trace between: *Anomalies / defects*, *Bandwidth statistics*, *Frame layer statistics*, *SLA statistics*, *BERT*, *Synchronization*.
7. Select the events to be monitored from the category you have selected in the previous step.

Once event logging is enabled and the monitored events have been selected, the equipment starts generating one trace file for each test started with *RUN* (or automatically through the Autostart/stop functionality).

5.4.2. Displaying Logs

You can either display trace files from finished tests or from the current test. You don't need to wait to the end of the test to display a trace file but logs are not upgraded in real time. You may need to re-open a trace file to display the results collected since the last time the file was opened. To display the trace files and browse the events they contain follow this procedure:

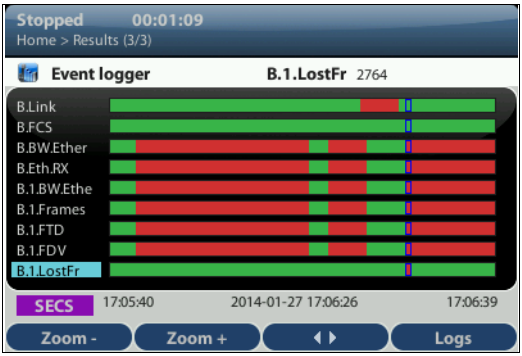

Figure 5.4: Ether.Genius / Ether.Sync / Ether.Giga event logger panel. This panel traces all events previusly selected from the event filter.

1. From the *Home* panel, go to *Results*,
   The *Results* panel is displayed.
2. Select *Event logger* to enter in event tracer.

3. Press *Logs* (*F4* contextual key).
   A list with all the available trace files is displayed. Files are identified by the measurement start date and time.
4. Select one trace fila and display it with the help of the *Open* (*F1* contextual key) control.
   Note: You can use *Delete* (*F2* contextual key) to delete one or various trace files from the list.
5. Use the cursors to browse the events registered at different times. You can set the display scale to seconds (SECS), minutes (MINS), hours (HOURS) and days (DAYS) with the help of the *Zoom+* (*F1* contextual key) and *Zoom-* (*F2* contextual key). Use the *F3* contextual key to switch between quick and slow event browsing.

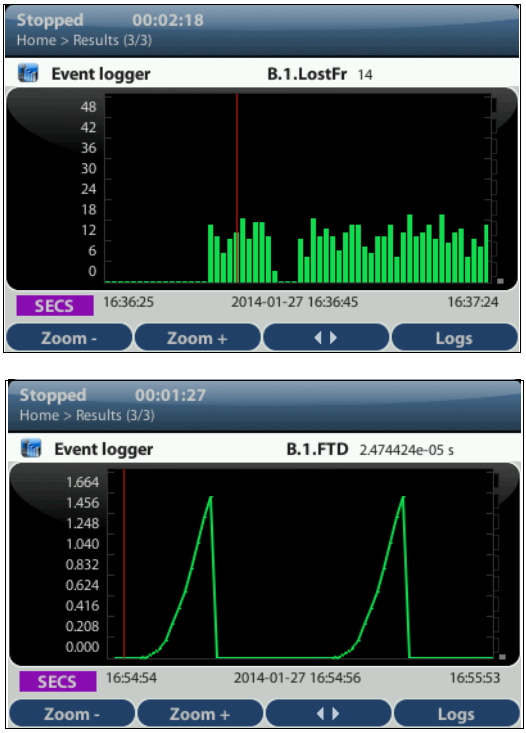

Figure 5.5: Ether.Genius / Ether.Sync / Ether.Giga detailed view of events: (a) Bar diagram (frame loss), (b) line diagram (frame delay).

6. Optionally, press ENTER to display detailed information about the currently selected event.
   Note: The controls described in the previous step are valid for the detailed view as well. Additionally, you can modify the vertical axis by pressing ENTER a second

time while you are in the detailed view. When you are in the vertical scale editing mode, you can use the cursors (*UP* an *DOWN*, changes the axis and *LEFT*, *RIGHT* changes the step size) to control axis variations. To confirm changes in the axis scale, press *ENTER* again.

5.5. BER Testing

Ether.Genius / Ether.Sync / Ether.Giga testers support Bit Error Rate (BER) testing over framed and unframed interfaces. The former computes the BER over Ethernet or IP interfaces, the later doesn't take into account the frame structure associated to the Ethernet interface and accounts for bit errors directly in the physical layer.

5.5.1. Framed BER Tests

Framed BER tests are compatible with Ethernet and IP interfaces. It is even possible to carry out a BER test between specific source and destination UDP ports.

Table 5.8: BERT Results

Metric	Description
TSE	Test Sequence Error. One TSE is equivalent to a single bit difference between the transmitted and the received test pattern (PRBS or other). This field is a cumulative counter of all the TSE events found from the beginning of the test. In framed Ethernet interfaces transmission errors may cause the test patterns to be altered and produce bit errors. However, the Ethernet FCS field contains a CRC-32 code designed for error detection. For this reason, the receiver (or any intermediate network element) may detect and discard the frame before any TSE is detected by the pattern analyser. The only way to account for TSE errors would be to recompute the FCS field after any transmission error.
BER	The Bit Error Ratio (BER) this is the ratio of the received TSE to the total amount of transmitted bits. The BER is one of the most fundamental quality parameters of TDM digital circuits. However, due to other degradation sources specific of statistically multiplexed networks (frame loss, variable frame delay), the performance description based on the BER may be incomplete in Ethernet / IP networks.

Table 5.8: BERT Results

Metric	Description
ES	This is the amount of Errored Second (ES) outcomes from the beginning of the test. An ES is defined as a second which contains at least one TSE or a higher order defect like an LSS
LSS	Loss of Sequence Synchronization. This event indicates that the expected test pattern does not match the actually received pattern. Frame loss events may cause temporary LSS.

Framed BER test is not as useful in Ethernet as it is in TDM networks but it can be used to trace connectivity with remote equipments and detect any temporary availability fault. The tester considers that a frame contains bit errors when the test pattern carried by the frame does not match the expected pattern but the checksum fields are still correct. For this reason, Ether.Genius / Ether.Sync / Ether.Giga Test Sequence Error (TSE) events are closer to corrupted frame events than to real transmission errors. In framed interfaces, transmission errors are likely to cause checksum errors and they are discarded before they can reach the test pattern analyser. In other words, for the tester, checksum errors (FCS errors, IPv4 errors, UDP errors) have higher precedence than bit errors. For this reason, transmission errors are usually accounted as a simultaneous checksum error and a lost frame event rather than a TSE event.

Ether.Genius / Ether.Sync / Ether.Giga testers include one single test pattern generator (Port A) and two pattern analysers (one for Port A and one for Port B). Pattern generator is bound with Flow 1 while traffic analysers are attached to Filter 1.

Before running a framed BER test, first you need to make sure that your equipment is configured in *Ethernet endpoint* or *IP endpoint modes* (See section 2.1). The traffic generator must be configured in the same way than any other traffic test, including the physical layer settings, frame settings, the network settings and the bandwidth profile (See chapter 4). The procedure to configure the test is as follows:

1. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
2. Select either *Port A* or *Port B* to enter in the port specific configuration.
3. Enter in the *Payload* menu.
4. Select *Flow 1* to enter in the flow 1 specific configuration.
   All settings related with payload configuration in the current flow are displayed.
5. Choose one of *BERT* in *Payload type* (See section 4.1.4).
6. Choose the bit pattern you are going to use for generation (Port A) and analysis (Port A and Port B) with the help of the *BERT patterns* control.

7. If you have configured *BERT patterns* to *User*, enter a 32-bit test pattern in *User payload* in hexadecimal format.

Once the test has been configured it can be started at any time with the help of the RUN button. BER results for the last (or current) test are available at any time:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Go to BERT
4. Select *Filter 1* to enter in the flow 1 specific results.
5. Check the *TSE BER*, *ES*, and *LSS* results.

5.5.2. Physical Layer BER Tests

The Physical Layer BER or L1 BER test measures the same parameters than the framed BER test but it is based on much simpler test patterns. Unlike the framed patterns, L1 patterns are fixed and not editable by the user. Some of them carry a minimum envelope like a preamble, an FCS or an IFG that enable some network elements to process the data stream in the same way than a framed flow. The advantage of L1 test patterns is that they are processed as a byte stream by the analyser (as opposed to a frame sequence). The result is a more accurate description of the transmission performance in terms of bit errors. On the other hand, L1 test patterns do not contain source and destination addresses or CoS labels to force the stream to follow an specific path in the network. Moreover, there is not a “bandwidth profile” you can use with L1 patterns, they are just byte sequences. At best, when they are carried within an Ethernet envelope, they use the whole capacity in the transmission interface.

Ether.Genius / Ether.Sync / Ether.Giga supports two types of L1 test patterns referred as L1-PCS-synchronized patterns and L2-compliant patterns. These are the properties of each type:

- *L1-PCS-synchronized patterns*: These patterns are defined by a sequence of *Physical Coding Sublayer* (PCS) symbols. For 1 Gb/s interfaces, L1-PCS-synchronized patterns are defined in terms of 8B/10B symbols. The pattern is not created as an 8-bit (8B) symbol and then converted to the correct 10-bit (10B) code. The result is that L1-PCS-synchronized patterns can be analysed at the bit level. L1-PCS-synchronized traffic passes through any circuit that has not any subsystem requiring valid MAC addressing or valid FCS codes.
- *L2-compliant patterns*: These test patterns are designed to resemble an Ethernet frame. The pattern is similar to a basic frame through de data link layer, including a preamble with a valid *Start of Frame Delimiter* (SFD), a FCS code and the correct encoding if the inter-frame gap. The pattern, however, overwrites all other parts of the frame, including MAC and IP addresses. L2-compliant patterns pass through any element looking for a valid Ethernet frame with FCS.

In Ether.Genius / Ether.Sync / Ether.Giga, L1-PCS-synchronized patterns are supported only by the optical interfaces. L2-compliant patterns are available from all physical interface configurations.

Table 5.9: Test Patterns for Unframed Operation

Setting	Description
RPAT	*Random Data Pattern*. This is a L1-PCS-synchronized test pattern defined in the NCITS TR-25-1999. This pattern is designed to provide energy across the entire frequency spectrum, and they provide a good basic BER test. The RPAT consists on the continuous repetition of the following sequence expressed in hexadecimal format: 0x3e-b0-5c-67-85-d3-17-2c-a8-56-d8-4b-b6-a6-65. The RPAT is available for optical L1 BER tests only.
JPAT	*Jitter Tolerance Pattern*. This is a L1-PCS-synchronized test pattern defined in the NCITS TR-25-1999. This test is designed for receiver jitter tolerance testing. The JPAT consists on the cyclic transmission of the following 8B/10B symbol sequence: D30.3 (repeated 192 times) and D21.5 (repeated 64 times)
SPAT	*Supply Noise test pattern*: This is a L1-PCS-synchronized test pattern defined in the NCITS TR-25-1999. It represents the worst-case power supply noise introduced by a transceiver. The SPAT consists on the cyclic transmission of the following hexadecimal pattern: 0xac-d4-ca-cd-4c (512 times). The SPAT is available for optical L1 BER tests only
HFPAT	*High Frequency test pattern*. This is a L1-PCS-synchronized test pattern defined in Annex 36A of the IEEE 802.3 standard. The purpose of the pattern is to test random jitter at a BER of 10^{-12}, and also to test the asymmetry of transition times. The HFPAT consists in the continuous repetition of the D21.5 code-group and it corresponds with the following bit sequence: 101010101010101010... The RPAT is available for optical L1 BER tests only.

Table 5.9: Test Patterns for Unframed Operation

Setting	Description
LFPAT	*Low Frequency test pattern*. This is a L1-PCS-synchronized test pattern defined in Annex 36A of the IEEE 802.3 standard. The purpose of this pattern is to test low-frequency random jitter and also to test PLL tracking error. The LFPAT consists on the continuous repetition of the K28.7 code-group and it corresponds with the following bit sequence: 11111000001111100000... The RPAT is available for optical L1 BER tests only.
MFPAT	*Mixed Frequency test pattern*.This is a L1-PCS-synchronized test pattern defined in Annex 36A of the IEEE 802.3 standard. The purpose of this pattern is to test the combination of random and deterministic jitter. The MFPAT consists on the continuous repetition of the K.28.5 code-group and it corresponds with the following bit sequence: 11111010110000010100111110101 1... The RPAT is available for optical L1 BER tests only.
LCRPAT	*Long Continuous Random test pattern*. This a L2-compliant test pattern defined in Annex 36A of the IEEE 802.3 standard. This pattern is designed to provide a broad spectral content and minimal peaking, allowing for the measurement of jitter at either the component or system level. The structure of the LCRPAT is based on cyclic transmission of a valid preamble, SDF sequence, the modified RPAT sequence repeated 126 times, a valid FCS and a 12-byte IFG The modified RPAT consists on the following sequence before the 8B/10B encoding: 0xbe-d7-23-47-6b-8f-b3-14-5e-fb-35-59.
SCRPAT	Sort Continuous Random test pattern. This is a L2-compliant test pattern defined in Annex 36A of the IEEE 802.3 standard. This pattern is designed to provide a broad spectral content and minimal peaking, allowing for the measurement of jitter at either the component or system level. The SCRPAT pattern structure is identical to the LCRPAT but the modified RPAT is repeated only 29 times resulting in shorter MAC frames.

Ether.Genius / Ether.Sync / Ether.Giga testers include one L1 test pattern generator (Port A) and two pattern analysers (one for Port A and one for Port B). For a L1 BER test it is required one test pattern generator and at least one analyser. Generators and analysers can be physically located in one or various units. Configuration of L1 BER tests is slightly different than the L2-L4 framed BER test. The procedure is as follows:

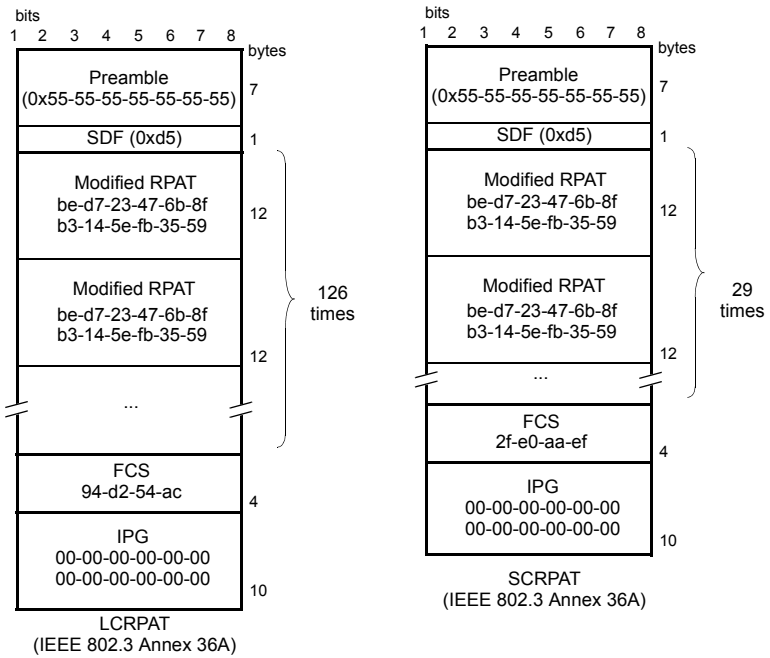

Figure 5.6: L2 compliant patterns supported by Ether.Genius / Ether.Sync / Ether.Giga.

1. Make sure that your tester is connected to the network.
   Note: Depending on your test setup you may need to connect various equipments, including traffic reflectors in the network.
   Note: Testing with L1-PCS-synchronized patterns (optical tests) does not require link establishment. Pattern generation is started as soon as testing starts without waiting for link auto-negotiation or other link establishment mechanism to finish.
1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Select *Mode* to enter in the mode selection menu
3. Choose *L1 Endpoint*. Confirm by pressing ENTER.
4. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.

5. Select either *Port A* or *Port B* to enter in the port specific configuration.
6. Go to L1 BERT pattern and select one or RPAT, JPAT, SPAT, HFPAT, LFPAT, MFPAT, LCRPAT or SCRPAT.
   Note: RPAT, JPAT, SPAT HFPAT, LFPAT and MFPAT patterns are not available in BER tests through electrical interfaces.
   Note: Depending on your particular test setup you may need to repeat this operation both Port A and Port B.

Once the test has been configured it can be started at any time with the help of the RUN button. BER results for the last (or current) test are available at any time:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Go to BERT
4. Select *Physical layer BERT* to enter in the L1 BER test results.
5. Check the *TSE BER*, *ES*, and *LSS* results.

Chapter 6
Multi-Stream Analysis

Ether.Genius / Ether.Sync / Ether.Giga are capable of processing and computing statistics over fractions of the Ethernet / IP traffic meeting specific conditions. The process of selecting a fraction of traffic is called filtering. The result of the filtering process is one or several traffic streams.

This chapter describes how to configure the tester for packet filtering and how to get statistics and results from each stream. These results can be classified in different categories:

- How many frames made up the traffic stream to be analysed. For example, filters can tell you how many frames have been received from an specific flow from the beginning of the test.
- The filter analysis tells what kind of data contains the traffic stream under analysis. For example, it the stream has at least some frames with the test payload required to measure SLA statistics you will be informed with an special indication.
- Filters are useful to get information about how much of the available bandwidth is being used by a traffic stream. If, for example, you choose to filter all broadcast IP packets (destination IPv4 address set to 255.255.255.255) you will measure how much traffic is consumed by broadcast traffic in your network like for example in ARP requests.
- Finally, the test unit provides per-filter information about critical SLA parameters like delay, delay variation or packet loss. It depends of the filtering criteria you are using that the SLA results you get are meaningful for your purposes or not.

6.1. Enabling and Disabling Filters

Traffic selection or filtering is configured by first enabling one or several filtering blocks and after that setting the filtering criteria. Ether.Genius / Ether.Sync / Ether.Giga support Ethernet, VLAN, IPv4 and TCP / UDP filters.

The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga is equipped with 8 filters. Each filter has a priority number. If one frame is selected by an specific filter it

will not be processed by any lower priority filters. In other words, traffic is processed by at most one matching filter, the one with the highest precedence.

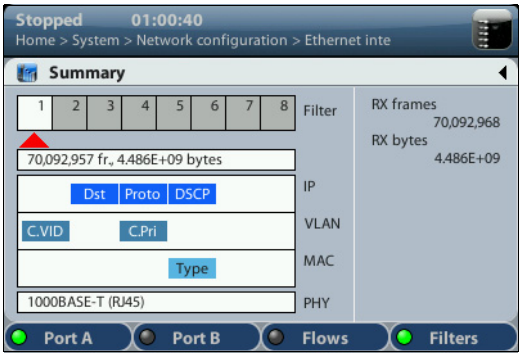

Figure 6.1: Filter summary panel. Filter status can be checked from this panel.

One filter admits three different configurations:

- *Block*: The filter blocks all the incoming traffic. No frame can match the filter. It can be considered that the filter remains disabled if it is configured to the block status.
- *Custom*: The filter accepts user defined matching rules like matching an specific VLAN, a group of source IP addresses or many others.
- *Match TX*: Matches the traffic configured in the corresponding flow from the traffic generator (Filter 1 matches Flow 1, Filter 2 matches Flow 2, etc.). Port B does not have traffic generation capabilities and therefore it attempts to match traffic transmitted at port A. The way traffic is matched in Port A and Port B filters is different. Port A expects that MAC addresses, IP addresses and UDP ports will be reversed when compared with the generator settings. For example, Port A filters look for the IP address configured as a source in the IP destination address field of incoming traffic. This address reversal rule does not apply to Port B.

Current filtering configuration is always available from the summary screen. In this screen use the F1 / F2 contextual buttons to choose between Port A or Port B filters and the F4 contextual button to display filtering configuration. The right / left cursors can be used to increase or decrease the filter index (*Filter 1*, *Filter 2*,...)

6.2. Configuring Filters

If a filter is left with its default configuration, it will not accept any frame. To allow the filter to accept and process frames, a correct filtering criteria must be configured before.

To configure the correct filtering criteria two decisions must be taken. First, it is necessary to know which frame fields are going to be matched and after that, which are the value or values to be matched. The first decision involves choosing whether the filtering is going to be done at MAC, IP or transport layer and which specific frame field or fields are going to be used for filtering (MAC addresses, IP addresses, Ethertype field, protocol or any other). The second is carried out by configuring the field value and sometimes a mask. The mask selects which field bits are taken into account when a frame is matched. Matching masks are not related to IP subnet masks even if they can be applied to IP addresses. Specifically, the binary representation of a matching mask does not need to be a sequence of '1' followed by a sequence of '0' like IP network masks are.

The generic procedure to configure one or several matching criteria with Ether.Genius / Ether.Sync / Ether.Giga is the following:

1. From the *Home* panel, go to *Setup*,
   The test port settings panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific configuration.
3. Enter in *Filters*
4. Select one of the filtering menus labelled as *Filter 1*, *Filter 2*, etc.
   All configuration items related with the filter configuration are displayed.
5. Configure Filter mode to one of *Block*, *Custom* or *Match TX* values
   Block and Match TX have fixed filtering rules but if you choose custom, different types of filtering rules are enabled for that filter: *Fixed offset*, *MAC*, C-*VLAN, S-VLAN, IPv4*, *UDP*.
6. If you have configured *Custom* filtering, select the matching rule.
7. For custom filters, choose a matching field and configure the matching mode for this field. Most of the matching fields have at least two matching modes. The *Equal* mode selects frames matching the configured value or values for the field and *Ignore* does not match any frame by the current field. Other matching modes may be available in specific fields.
8. Configure the field value to be matched by the filter.
9. If the matching field has this capability, enter the mask value. To select a single value, set of the mask bit values to all ones.
10. Optionally, configure more matching rules for the current filter by repeating steps 3, 4, 5, 6 and 7 as many times as necessary.

6.2.1. MAC Selection

MAC frames are envelopes in which the Ethernet frames are sent and received. MAC frame format is currently specified by the standard IEEE 802.3. This format is shared by all existing Ethernet interfaces thus making Ethernet the most scalable transmission technology currently available.

The ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga provide frame selection based on the MAC address source and destination and Ethertype value. It is possible

to configure a matching mask for all three fields to select a value set rather than a single value.

Table 6.1: MAC Selection

Field	Description
Source MAC address match	Enables selection by source MAC address in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The source MAC address is ignored and not taken into account for the purpose of the filter. No frame is selected by the filter when *Ignore* has been configured. • *Equal to*: All frames matching a user configurable pattern are allowed to pass through the filter. The matching pattern is configured with the help of the *Source MAC address* and *Source MAC address mask* fields.
Source MAC address	This is a 48-bit MAC address in the standard hexadecimal-digit format *XX:XX:XX:XX:XX:XX*. Source addresses matching some or all bits of the address configured here will be allowed to pass through the filter. Bits used for matching frames to this filter are configured in the *Source MAC address mask*.
Source MAC address mask	This is the mask for the source MAC address filter selection rule. Before comparing the *Source MAC address* field with the frame addresses, bit wise AND operations are carried out between the value configured here and the *Source MAC address* field so that only the values surviving the AND are taken into account for matching the filter.
Destination MAC address match	Enables selection by destination MAC address in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source MAC address selection* setting.
Destination MAC address	This is a 48-bit MAC address in the standard hexadecimal-digit format *XX:XX:XX:XX:XX:XX*. Destination addresses matching some or all bits of the address configured here will be allowed to pass through the filter. Bits used for matching frames to this filter are configured in the *Destination MAC address mask*.

Table 6.1: MAC Selection

Field	Description
Destination MAC address mask	This is the mask for the destination MAC address filter selection rule. Before comparing the *Destination address* field with the frame addresses, bit wise AND operations are carried out between the value configured here and the *Destination address* field so that only the values surviving the AND are taken into account for matching the filter.
Ethertype match	Enables selection by Ethertype value in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source address selection* and *Destination address selection* settings.
Ethertype	This setting contains a 2-byte field that constitutes the Ethertype value to be matched in the incoming traffic. Ethertypes matching some or all bytes of the value configured here will be allowed to pass through the filter. Bits used for matching frames to this filter are configured by means the *Ethertype mask*.
Ethertype mask	This is the mask for the Ethertype filter selection rule. Before comparing the *Ethertype* field with the frame Ethertype, bit wise AND operations are carried out between the value configured here and the *Ethertype* field so that only the values surviving the AND are taken into account for matching the filter.

6.2.2. C-VLAN and S-VLAN Selection

Within enterprise networks, VLANs are important because they enable network segmentation on an organisational basis, by functions, project teams or applications, rather than on a physical or a geographical basis. The network can be reconfigured through software, instead of physically unplugging and moving devices or wires.

VLANs are an important contribution to scalable Ethernet networks, because they limit broadcast traffic inherent to the bridging mechanism. Large amounts of broadcast traffic may damage performance and even collapse network equipment, which is why it must be controlled.

Standard IEEE 802.1Q specifies the most popular VLAN frame format. VLAN frames carry a 16-bit header which specifies the VLAN Identifier (VID) and the frame priority within the VLAN. Many carrier Ethernet networks use the VID for segmentation just like enterprises. The VID in carrier Ethernet networks is used by service providers as general purpose identifier. They can be associated to an specific service, customer,

node or several of them at the same time. Sometimes, service providers use a two-level VLAN structure. Levels are designated as customer VLAN (C-VLAN) and service VLAN (S-VLAN). This two-structure is know as Q-in-Q. The standardised version of the Q-in-Q frame is defined in IEEE 802.1ad. Ether.Genius / Ether.Sync / Ether.Giga testers support both single-level Q-frame and the two-lever Q-in-Q-frame.

Table 6.2: C-VLAN and S-VLAN Selection

Field	Description
VID match	Enables selection by VID in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The VID is ignored and not taken into account for the purpose of the filter. No frame is selected by the filter when *Ignore* has been configured. • *Equal to*: All frames matching a user configurable pattern are allowed to pass through the filter. The matching pattern is configured with the help of the *VID* field.
VID	This setting contains a 10-bit identifier that constitutes the VID value to be matched in the incoming traffic. It is possible to match the S-VID or the C-VID through separated entries in the filtering menu. For single-tagged frames it is assumed that the frame does not contain S-VID field and therefore configuration is done through the C-VLAN menu.
Priority codepoint match	Enables selection by IEEE 802.1Q/p priority bits in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *VID match* field.
Priority codepoint	This setting contains a 3-bit identifier that constitutes the priority value to be matched in the incoming traffic. It is possible to match the S-VLAN or the C-VLAN priority codepoints through separated entries in the filtering menu. For single-tagged frames it is assumed that the frame does not contain S-VLAN priority codepoint field and therefore configuration is done through the C-VLAN menu.
Drop-eligible indicator match	Enables selection by the S-VLAN drop-eligible indicator in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and with ones are blocked. The available configuration values for this field are the same that for the *VID match* field.

Table 6.2: C-VLAN and S-VLAN Selection

Field	Description
Drop-eligible indicator	This is a single bit field that is used to signal which frames have precedence when the node has to drop some information due to congestion or other causes. The Drop-eligible indicator is defined only for the S-VLAN tag of double-tagged frames.For this reason is only available for S-VLAN matching rules.

6.2.3. MPLS Selection

MPLS traffic carry one or various 4-byte headers. Each of these headers is made up of a 20-bit MPLS label, a 3-bit traffic class identifier and other fields (See section 4.2.2). The filtering capabilities of Ether.Genius / Ether.Sync / Ether.Giga can be used to match MPLS header field of IP packets carrying up to two different labels.

Table 6.3: MPLS Selection

Field	Description
Bottom label match	Enables selection based on the bottom MPLS label in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The bottom MPLS label is ignored and not taken into account for the purpose of the filter. No frame is selected by the filter when *Ignore* has been configured. • *Equal to*: All frames matching a user configurable MPLS label are allowed to pass through the filter. This label is configured with the help of the *Bottom label* field.
Bottom label	Bottom of the stack MPLS label. It is 20-bit MPLS label in decimal format. Labels are used for switching packets in MPLS networks. Labels have local meaning only. That means that a single LSP could have different MPLS labels in different links between different routers

Table 6.3: MPLS Selection

Field	Description
Bottom traffic class match	Enables selection based on the class of service bits included in the top MPLS label in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The bottom class of service field is ignored and not taken into account for the purpose of the filter. No frame is selected by the filter when *Ignore* has been configured. • *Equal to*: All frames matching a user configurable class of service field are allowed to pass through the filter. The class of service field is configured with the help of the *Bottom traffic class* field.
Bottom traffic class	Bottom 3-bit MPLS CoS identifier. It was first thought that this field could carry the 3 IPv4 Type-of-Service (ToS) bits, but currently, the ToS field is being replaced by 6-bit *Differentiated Services Code Points* (DSCP). This means that only a partial mapping of all the possible DSCPs into this field is possible.
Top label match	Enables selection based on the top MPLS label in the current filter. Top label matching makes sense only if the received packets carry two MPLS labels. In this case the top label is the one placed closer to the layer two header. The available configuration values for this field are the same that for the *Bottom label match* setting.
Top label	Top of the stack MPLS label. It is 20-bit MPLS label in decimal format. Labels are used for switching packets in MPLS networks. Labels have local meaning only. That means that a single LSP could have different MPLS labels in different links between different routers
Top traffic class match	Enables selection based on the class of service bits included in the top MPLS label in the current filter. The available configuration values for this field are the same that for the *Top label match* setting.
Top traffic class	Top 3-bit MPLS CoS identifier. It was first thought that this field could carry the 3 IPv4 Type-of-Service (ToS) bits, but currently, the ToS field is being replaced by 6-bit *Differentiated Services Code Points* (DSCP). This means that only a partial mapping of all the possible DSCPs into this field is possible.

6.2.4. IPv4 Selection

Ether.Genius / Ether.Sync / Ether.Giga filtering capabilities can be programmed to match fields within the IPv4 datagram. It is currently supported IP datagram matching based on source IP address, destination IP address, protocol and DSCP. Source and destination IP addresses can be matched by means selection masks.

Table 6.4: IPv4 Selection

Field	Description
Source IPv4 address match	Enables selection by source IPv4 address in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The source IP address is ignored and not taken into account for the purpose of the filter. No frame is selected by the filter when *Ignore* has been configured. • *Equal to*: All frames matching a user configurable pattern are allowed to pass through the filter. The matching pattern is configured with the help of the *Source IPv4 address* and *Source IPv4 address mask* fields.
Source IPv4 address	This is a 32-bit IPv4 address in the standard four-dotted decimal format *A.B.C.D*. Source addresses matching some or all bits of the address configured here will be allowed to pass through the filter. Bits used for matching frames to this filter are configured in the *Source address mask*.
Source IPv4 address mask	This is the mask for the source IPv4 address filter selection rule. Before comparing the *Source address* field with the frame addresses, bit wise AND operations are carried out between the value configured here and the *Source address* field so that only the values surviving the AND are taken into account for matching the filter.
Destination IPv4 address match	Enables selection by destination IPv4 address in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv4 address match* setting.
Destination IPv4 address	This is a 32-bit IPv4 address in the standard four-dotted decimal format *A.B.C.D*. Destination addresses matching some or all bits of the address configured here will be allowed to pass through the filter. Bits used for matching frames to this filter are configured in the *Destination address mask*.

Table 6.4: IPv4 Selection

Field	Description
Destination IPv4 address mask	This is the mask for the source IPv4 address filter selection criteria. Before comparing the *Source address* field with the frame addresses, bit wise AND operations are carried out between the value configured here and the *Source address* field so that only the values surviving the AND are taken into account for matching the filter.
IP protocol match	Enables selection by the 1-byte IPv4 protocol field in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv4 address match* setting.
IP protocol	Configures the protocol to be filtered when protocol selection is enabled. The available configuration values are the following: • *Numeric*: Use this control if the traffic to be matched is different of UDP, TCP and ICMP and it has an specific protocol identifier assigned by the IANA. • *UDP*: Matches traffic with a *User Datagram Protocol (UDP)* envelope. Traffic commonly transported over UDP includes IP voice, IP video and DNS. • *TCP*: Matches traffic carried over the Transfer Control Protocol (TCP). Most data applications (web, file transfer, e-mail...) ere normally based on TCP transport. • *ICMP*: Matches *Internet Control Message Protocol* packets. IP operation and maintenance traffic like ping use ICMP.
IP protocol number	This setting contains an 8-bit word that constitutes the protocol identifier to be matched in the incoming traffic. Configuring this field to 17 is equivalent of setting the *Standard protocol selection* to UDP. TCP uses 6 as the protocol number and ICMP uses number 1. This control is enabled only if Standard protocol selection has been previously set to *Numeric*.
DSCP match	Enables selection by the 6-bit DSCP field in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv4 address match* field.

Table 6.4: IPv4 Selection

Field	Description
DSCP	This setting contains a 6-bit word in decimal format that constitutes the DSCP to be matched in the incoming traffic.

6.2.5. IPv6 Selection

Version 6 of the IP protocol is increasingly important in current network deployments. For this reason, Ether.Genius / Ether.Sync / Ether.Giga supports filtering based on various IPv6 packet fields including addresses, CoS marks, flow labels and higher layer protocol identifiers.

Table 6.5: IPv6 Selection

Field	Description
Source IPv6 address match	Enables selection by source IPv6 address in the current filter. The value configured here is used to determine which packets are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The source IP address is ignored and not taken into account for the purpose of the filter. All packets are allowed to pass through the filter when *Ignore* is configured if they are not blocked by other selection rule. • *Equal to*: All packets matching a user configurable pattern are allowed to pass through the filter. The matching pattern is configured with the help of the *Source IPv6 address* and *Source IPv6 address mask* objects.
Source IPv6 address	This is a 128-bit IPv6 address in the A:B:C:D:E:F:G:H format, where A, B, C, D, E, F, G and H are hexadecimal numbers between *0000* and *ffff*. Source addresses matching some or all bits of the address configured here will be allowed to pass through the filter. Bits used for matching packets to this filter are configured by means the *Source IPv6 address mask* field.
Source IPv6 address mask	This is the mask for the source IPv6 address filter selection rule. Before comparing the *Source IPv6 Address Match* field with the actual IPv6 addresses, bit wise *AND* operations are carried out between the value configured here and the source address so that only the values surviving the *AND* are taken into account for matching addresses.

Table 6.5: IPv6 Selection

Field	Description
Destination IPv6 address match	Enables selection by destination IPv6 address in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv6 address match* setting.
Destination IPv6 address	This is a 128-bit IPv6 address in the A:B:C:D:E:F:G:H format, where A, B, C, D, E, F, G and H are hexadecimal numbers between *0000* and *ffff*. Destination addresses matching some or all bits of the address configured here will be allowed to pass through the filter. Bits used for matching packets to this filter are configured by means the *Destination IPv6 address mask* field.
Destination IPv6 address mask	This is the mask for the destination IPv6 address filter selection rule. Before comparing the *Destination IPv6 Address Match* field with the actual IPv6 addresses, bit wise *AND* operations are carried out between the value configured here and the source address so that only the values surviving the *AND* are taken into account for matching addresses.
Next Header match	Enables selection by the 8-bit IPv6 next header field in the current filter. The value configured here is used to determine which packets are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv6 address match* setting.
Next Header	Configures the protocol identifier to be filtered when *Next Header matching* is enabled. The available configuration values are the following: • *Numeric*: Use this control if the traffic to be matched is different of UDP, TCP and ICMP and it has an specific protocol identifier assigned by the IANA. • *UDP*: Matches traffic with a *User Datagram Protocol (UDP)* envelope. Traffic commonly transported over UDP includes IP voice, IP video and DNS. • *TCP*: Matches traffic carried over the Transfer Control Protocol (TCP). Most data applications (web, file transfer, e-mail...) ere normally based on TCP transport. • *ICMP*: Matches *Internet Control Message Protocol* packets. IP operation and maintenance traffic like ping use ICMP.

Table 6.5: IPv6 Selection

Field	Description
Next Header number	This object contains an 8-bit word that constitutes the next header identifier to be matched in the incoming traffic. For example, configuring this object to 17 matches UDP traffic, TCP uses 6 as the protocol number and ICMPv6 uses number 58. The *Next Header number* configuration field is enabled only if *Next Header* is configured to *Numeric*.
Flow Label match	Enables selection by the 20-bit IPv6 flow label field in the current filter. The value configured here is used to determine which packets are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv6 address match* setting.
Flow label	The *Flow Label* IPv6 field contains a 20-bit word that identifies an unidirectional data flow. These labels remain at disposal of intermediate routers for stateful and stateless processing at flow level. For example, the flow label could be used to prevent load balancing on a particular traffic flow.
DSCP match	Enables selection by the 6-bit differentiated services code point (DSCP) field in the current filter. The value configured here is used to determine which packets are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source IPv6 address match* field.
DSCP	This object contains a 6-bit word in decimal format that constitutes the DSCP to be matched in the incoming traffic.

6.2.6. UDP Selection

Some applications do not require reliable transmission at the transport layer either because they implement their own error control mechanisms or because the mechanisms used by TCP are too slow for them. These applications can use the light

weight User Datagram Protocol (UDP). Like TCP, UDP provides communications through ports to applications but it doesn't have any error recovery capability.

Table 6.6: UDP Selection

Field	Description
Source port match	Enables selection by source UCP port in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: The source port is ignored and not taken into account for the purpose of the filter. No frame is selected by the filter when *Ignore* has been configured. • *Range*: All frames with a destination port in an specified range are allowed to pass through the filter. The port range is specified with the help of the *Minimum source port* and *Maximum source port* fields.
Minimum source port	This is the minimum 16-bit UDP source port allowed to pass through the *Source port match* filter. The port is configured and displayed in decimal format.
Maximum source port	This is the maximum 16-bit UDP source port allowed to pass through the *Source port match* filter. The port is configured and displayed in decimal format.
Destination port match	Enables selection by the 2-byte UDP protocol field in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are the same that for the *Source address match* field.
Minimum destination port	This is the minimum 16-bit UDP destination port allowed to pass through the Destination port selection filter. The port is configured and displayed in decimal format.
Maximum destination port	This is the maximum 16-bit UDP destination port allowed to pass through the Destination port selection filter. The port is configured and displayed in decimal format.

6.2.7. Fixed Offset Selection

Generic selection is the matching mode to be used when the Ethernet frames carry uncommon protocols or if inspection beyond the UDP and TCP transport protocols is

required. This selection mode defines an offset and a mask. Frames matching the specified mask in the configured offset are selected.

Table 6.7: Generic Selection

Field	Description
Filter match	Enables fixed offset selection in the current filter. The value configured here is used to determine which frames are allowed to pass through the filter and which ones are blocked. The available configuration values for this field are: • *Ignore*: Generic matching is ignored and not taken into account for the purpose of the filter. All frames are allowed to pass through the filter when *Ignore* is configured if they are not blocked by other selection criteria. • *Equal to*: All frames matching a user configurable pattern are allowed to pass through the filter. The matching pattern is configured with the help of the *Payload selection*, *Offset (bytes)*, *Match code* and *Mask* fields.
Payload selection	Defines the payload type and the reference point within the frame to start counting the filter offset. The reference can be the beginning of the MAC, IPv4 or UDP payload depending on the chosen value. It is also possible to set the reference to the beginning of the Ethernet frame (first byte immediately after the SDF) by configuring *Whole frame* in this field. The *Payload selection* field is shared by all the port A or port B filters. If the *Frame start* field is modified for one specific filter, the remaining filters of the same port will be automatically configured to the same value.
Offset (bytes)	This field defines the offset expressed in bytes from the reference point defined with the *Payload selection* control. The value 0 corresponds with the first byte of the MAC, IPv4 or UDP payload (or the first frame byte, if *Payload selection* is set to *Whole frame*). If due to the limited frame size, some or all the byte positions defined by the *Offset (bytes)* field, do not exist in the corresponding payload, the equipment will consider that the frame does not match the filtering criteria. The offset field is shared by all the port A or port B filters. If the *Offset (bytes)* field is modified for one specific filter, the remaining filters of the same port will be set to the same value.

Table 6.7: Generic Selection

Field	Description
Match code	16-bit code expressed with four hexadecimal digits used to match frames.in the current filter.
Mask	This is a mask for the generic filter match code. Before comparing the *Match code* with the selected bytes in the Ethernet frame, bit wise AND operations are carried out between the value configured here and the *Match code* field so that only the values surviving the AND are taken into account for matching the filter.

6.3. Per-Stream Counts and Statistics

Once the filter has been configured, it is usually desirable to know how many matching frames have been found for each filter and what are the matching traffic properties. Ether.Genius / Ether.Sync / Ether.Giga testers offer a complete set of statistics for each of the eight filtered streams, including bandwidth and SLA metrics. To get general statistics about filtered frames follow this procedure:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Go to *Filters*
   A table with dedicated information for each filter is displayed.
4. Check the *Frames*, *Bytes* and *Traffic* results.

Table 6.8: General Filter statistics

Field	Description
Frames	Count of all frames matching the selection rule for the current filter.
Bytes	Byte count corresponding to the traffic stream associated with the current filter. Note that only with the *Frames* and *Bytes* results it is possible to compute new statistics like the average frame size of the matching traffic.

Table 6.8: General Filter statistics

Field	Description
Traffic	Displays information about the traffic type detected during the last second. There are three possible results: • *None*: No traffic has been detected matching the filtering rules for the current filter during the last second. Either there is no traffic received in the port or the filter rules are too restricting to match any frame. • *Other*: At least one frame of network traffic matching the filter rules has been found during the last second. Traffic may come from anywhere in the network. • *SLA*: At least one frame carrying the ALBEDO Telecom SLA payload and matching the filter rules has been found during the last second. The SLA payload is ALBEDO proprietary and therefore it has probably been transmitted by an ALBEDO Telecom equipment. SLA statistics for the current filter cannot be calculated if no SLA traffic matching the filtering rules is received.

6.3.1. Bandwidth Statistics

Filter specific bandwidth statistics have the same meaning that the port-wide bandwidth statistics (See section 5.1.4) but in this case they are restricted to the traffic matching the filtering rules rather than the whole traffic received in the test port. Let's suppose that you want to know how much bandwidth is using an IPTV stream in your network. If you know some information about the stream like for example the destination UDP port used by the stream packets or the destination multicast IP address you can use them as a filtering rules and the bandwidth statistics for the filtered traffic will provide the result in the Ethernet, IP and UDP layers.

The procedure to display the bandwidth statistics for each of the configured filters is as follows:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Go to *Bandwidth statistics.*
4. Select one of the filter menus labelled as *Filter 1*, *Filter 2*, etc.
   All the result menus related with the current filter are displayed.
5. Check the *Eth. (current)*, *Eth. (min.)*, *Eth (max.)*, *IPv4 (current)*, *UDP (current)* results.

6.3.2. SLA Statistics

SLA statistics are more similar to BER results than bandwidth statistics or error counts: Like it happens with BER results, SLA statistics require an special frame / packet payload. In this case, the payload to be used is the SLA payload (See section 4.1.4, See section 4.4.6). That means that the tester you are using as a traffic generator must be configured to use this payload to obtain any valid result.

It is possible to use the same test port for SLA traffic generation and analysis (Port A) or use Port A for traffic generation and Port B for analysis. Using Port A for generation and analysis requires a loopback device somewhere in the network to redirect the traffic towards the origin (See section 2.2.3). It is also possible to generate traffic with one tester and use a remote equipment to analyse this traffic but delay statistics will be probably wrong or at least will have a poor accuracy due to the lack of a common timing source for the generator and the analyser.

Ether.Genius / Ether.Sync / Ether.Giga support SLA tests both in bridged environments (*Ethernet endpoint* test mode) and routed environments (*IP endpoint* test mode). Traffic generation in pass through mode is not supported, for this reason, there are no SLA results in *Ethernet / IP Through* mode.

The testers measure and represent up to eight sets of SLA statistics for each of the eight filter blocks. An important advise is that strictly, SLA results are meaningful only for frames carrying the SLA payload. Other traffic accounted in the same filter may be subject to different delay conditions. However, if one filter is receiving mixed (SLA and not SLA) traffic, it is possible to guarantee the representativeness of the SLA results if the correct filtering conditions are applied for the traffic. For example if it is know that the DUT / SUT is applying different forwarding policies depending only on the DSCP field, filters should be configured to use the DSCP as a filtering rule.

SLA results are timed, you need to start a test with run to enable the equipment to collect results. If you are using the equipment also as a test traffic generator you will also need to start a test to enable traffic generation. Once the test is configured and running you can access to the results. To configure and run an SLA test follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Go to *SLA statistics*.
4. Select one of the filtering menus labelled as *Filter 1*, *Filter 2*, etc.
   All the result menus related with the current filter are displayed.
5. Check the *FTD. (current)*, *FTD. (maximum)*, *FTD (minimum)*, *FTD (average), FTD (standard dev.)*, *FTD (range)*, *FDV (current)*, *FDV (maximum)*, *FDV (average)*,

Lost frames, *FLR*, *SES*, *PEU (%)*, *Out-of-sequence frames* and *Duplicated frames* results.

Table 6.9: SLA Statistics

Field	Description
FTD (current)	Current value of the point-to-point Ethernet Frame Delay computed as specified in ITU-T Y.1563 and expressed in ms. To display the current FTD it is required a previous initialization of a test with the RUN button.
FTD (maximum)	Peak value of FTD registered from the beginning of the test. The *FTD (max)* value is a packet by packet computed statistic. All relevant SLA packets are considered for calculation of this statistics.
FTD (minimum)	Minimum value of FTD registered from the beginning of the test. The *FTD (min)* value is a packet by packet computed statistic. All relevant SLA packets are considered for calculation of this statistics.
FTD (average)	Average of all registered FTD values from the beginning of the test. All relevant SLA packets are considered for calculation of this statistics (no sampling process is involved in the calculation of this statistic).
FTD (standard dev.	Standard deviation (positive square root of the variance) corresponding with all the registered FTD values from the beginning of the test. The standard deviation quantifies the delay dispersion or delay variation found in the DUT / SUT.
FTD (range)	This performance metric is the result of subtracting *FTD (min)* from *FTD (max).* The range is an alternative way to the standard deviation to quantify the delay dispersion found in the DUT / SUT.
FDV (current)	Current value of the jitter computed as per RFC 3393 and RFC 1889. Delay variation is computed over consecutively transmitted packets. The jitter value is smoothed with the function defined in RFC 1889 before being displayed. The current FDV constitutes still another metric to quantify the delay variation in the DUT / SUT. To display the current jitter it is required a previous initialization of a test with the RUN button.

Table 6.9: SLA Statistics

Field	Description
FDV (maximum)	Smoothed jitter maximum computed from the beginning of the test. The *FDV (max)* is a packet-by-packet statistic that does not involve any sampling process.
FDV (average)	Average of all individual delay variation values computed from the beginning of the test. Each individual delay variation is evaluated as the absolute value of the FTD associated to a given frame minus the FTD associated to the frame transmitted next. All possible consecutive frame transmission events are taken into account for the calculation of this performance metric. The only exception to this rule is if one or both frames are lost.
Lost frames	Total amount of lost frames from the beginning of the test. One frame is considered to be lost if it is never received or if it is received with a delay larger than 10 seconds.
FLR	Ratio of the total amount of lost frames to the total transmitted frames from the beginning of the test. Definition of the FLR parameter follows ITU-T Y.1563.
SES	This is the amount of Severely Errored Second (SES) outcomes from the beginning of the test. The SES is computed as specified in ITU-T Y.1563. The frame loss threshold used to declare a SES is 50% of the frames that made up the transmitted block.
PEU (%)	Percent Ethernet service Unavailability defined as per ITU-T Y.1563 and recorded from the beginning of the test. The PEU constitutes a availability performance figure that informs about the percentage of time that the DUT / SUT has been not available for transmit / receive data.
Out-of-sequence frames	Total amount of packets received with an unexpected sequence number. Duplicated sequence numbers are not taken into account for computing the Out-of-sequence frame statistic. This metric is based on the definitions given in RFC 5236..

Table 6.9: SLA Statistics

Field	Description
Duplicated frames	Total amount of frames received with a duplicated sequence number. A triplicated frame event is accounted as two duplicated frame events. The same reasoning is applied for frames repeated more than three times. The duplicated frames metric is based on the definitions given in RFC 5236.

Chapter 7
Automatic Performance Tests

Automatic tests are different to the measurements explained in previous chapters in that they are usually easier to configure and run. Specifically, the user does not need to worry about which bandwidth profile to use or which test payload to configure. However, users are still required to enter the correct MAC and IP addresses, CoS marks, and other frame and network configuration through the menu system (See section 4.1.2 and See section 4.4.4). Due to the way automatic tests use the bandwidth profile settings, the equipment may need to generate large amounts of traffic. This traffic may cause congestion in some unprepared networks and damage performance of any service already deployed. For this reason, users are advised to use automatic measurements with care.

The second relevant property of automatic tests is that they provide a clear pass or fail result that is easier to understand than a numeric latency figure or a bandwidth statistic. Thresholds for the pass / fail results can be tuned through specific menus.

ALBEDO Telecom Ether.Genius / Ether.Sync / Ether.Giga currently support two different automatic measurements: The IETF RFC 2544 tests and the Ethernet service activation test methodology (eSAM) based on the ITU-T Y.1564 standard. The former has been used for many years and its a very well established network benchmarking mechanism. The latter has been introduced more recently but it has several advantages over the RFC 2544 tests like faster execution, delay variation result, support for multiservice environments and support for coloured traffic.

7.1. Performance Assessment with the RFC 2544 Test

The RFC 2544 is an IETF standard that describes benchmarking tests for network devices. Vendors can use these tests to measure and outline the performance characteristics of their Ethernet and IP switching equipment. As these tests follow standard procedures, they also make it easier for customers to make sense of the glitzy marketing-speak employed by most vendors.

The tests described in the document aim to evaluate how a device would act in a real situation. The RFC 2544 describes out-of-service tests, which means that real traffic

must be stopped and the tester will generate specific frames to evaluate throughput, latency, frame loss rate, burst tolerance, overload conditions and recovery.

To configure an RFC 2544 under Ether.Genius / Ether.Sync / Ether.Giga follow these steps:

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
3. Choose between *Ethernet endpoint* or *IP endpoint* with the *Mode* setting.
4. Depending on your test setup (See section 2.2.3), configure *Test method* to *One-way (A > B)* or *Two-way (A > A)*.
5. Configure *Performance test* to *RFC 2544*.
6. Go to *RFC 2544*.
   The panel with specific configuration of the RFC 2544 throughput, latency, frame loss, back-to-back frames and system recovery time tests is displayed.
7. Configure the test frame sizes with the help of the *Frame sizes* multiple selection list. The available frame sizes are: *64*, *128*, *256*, *512*, *1024*, *1280*, *1518* and *User-defined*.
8. If *User-defined* is enabled in the previous step, configure the user frame size to something between 64 and 10000 bytes through the *User frame size* control.
9. Configure the *Throughput test* (See section 7.1.1), *Latency test* (See section 7.1.2), *Frame Loss test* (See section 7.1.3), *Back-to-back test* (See section 7.1.4) and *System recovery test* (See section 7.1.5).
10. Set the pass / fail thresholds for the throughput, latency, frame loss, back-to-back frames and system recovery tests from the *Performance objectives* menu (See section 7.1.6)
11. Configure the *Frame layer* (See section 4.1.2) in Port A for Flow 1. Parameters to be configured are source and destination MAC addresses, VLANs, etc.
    Note: All flows different to *Flow 1* are disabled in RFC 2544 tests. This test is only compatible with a single flow transmission.
12. If you are working in *IP endpoint* mode, configure the *Local profile* (See section 2.3) and the *Network layer* (See section 4.4.4). Parameters to be configured are IP addresses, DSCPs, etc.

7.1.1. Throughput

The aim of a *throughput test* is to determine the maximum number of frames per second that the device can process and forward without dropping or losing any. Put in simple terms, the procedure used to compute the throughput is as follows:

1. Send a certain number of frames at a specific rate through the DUT / SUT and count the frames transmitted by the DUT / SUT.

2. If the count of transmitted frames is equal to the count of received frames, increase the transmission rate and re-run the test. Otherwise, reduce the transmission rate to be used in the next trial.
3. Re-run the test until fewer frames are transmitted than received by the DUT / SUT. The throughput is the fastest rate at which the count of test frames transmitted by the DUT is equal to the number of test frames sent to it by the measurement equipment.

Table 7.1: RFC 2544 Throughput Test Settings

Setting	Description
Enable	Enables or disables the RFC 2544 throughput test. The throughput test will not run with the remaining RFC 2544 tests if they are not previously enabled. Some RFC 2544 tests require the throughput results to be executed. If the throughput test is not executed, the throughput results are replaced by the nominal capacity of the transmission medium (1 Gb/s, 100 Mb/s or 10 Mb/s, depending on configuration and transmission medium).
Maximum rate (%)	It is the maximum bit rate the DUT / SUT supports due policing or other bandwidth limiting mechanisms. The *Maximum rate (%)* constitutes the upper level of the throughput result. The *Maximum rate (%)* is configured as a percentage of the nominal transmission rate configured through link auto-negotiation or fixed by the network administrators. If the DUT / SUT is not using bandwidth limiting mechanisms or the maximum rate is unknown, configure this field to 100 %. Test results should not change but it will take more time to get the same results with the same accuracy level.
Resolution (%)	Minimum throughput, expressed as a percentage of the nominal channel capacity, that can be distinguished by the RFC 2544 throughput algorithm. This algorithm finishes when the distance between the real throughput and the value estimated by the algorithm is smaller than the resolution. Configuring an smaller value of the resolution increases measurement accuracy but also increases the time needed to finish the throughput test.

Table 7.1: RFC 2544 Throughput Test Settings

Setting	Description
Trial duration (s.)	This value corresponds with the duration of a single throughput trial test expressed in seconds. Each algorithm test cycle contains one trial. The equipment transmits test traffic during the time specified by this setting before deciding by means the analysis of the number of lost frames whether to continue or not. Some equipments may need some time before they start dropping frames. If the trial duration is too short this may cause the estimated throughput to be larger than it will in stationary transmission conditions.
Max. frame loss (%)	The equipment decides whether it is operating above or below the actual DUT / SUT throughput by comparing the number of transmitted and received frames. It decides that the current transmission rate is higher than the real throughput if the frame loss ratio is larger than the *Max.frame loss* (%) configured here. In the same way, the tester decides that the current test rate is smaller than the actual throughput if the computed frame loss ratio is smaller than the *Max. frame loss (%)*.

7.1.2. Latency

This test determines the latency inherent in the DUT / SUT. The initial data rate is based on the results of the throughput test. Typically, packets are timestamped using the SLA payload (See section 4.1.4, See section 4.4.6), and the time they take to travel through the DUT / SUT is measured.

In order to determine the latency usually you first measure the throughput for the DUT / SUT at each of the defined frame sizes, and send a stream of frames through the DUT / SUT at the determined throughput rate to a specific destination. If you don't measure throughput the latency will be measured at the nominal capacity for the transmission medium.

The time at which this frame is completely transmitted is recorded, and this will be timestamp A. The receiver of the test equipment must recognize the tag information in the frame stream and record the reception time of the tagged frame. This will be

timestamp B. The latency is the difference between timestamp B and timestamp A, according to the definition found in RFC 1242.

Table 7.2: RFC 2544 Latency Test Settings

Setting	Description
Enable	Enables or disables the RFC 2544 latency test. The latency test will not be executed if is not previously enabled by this control.
Trial duration (s.)	This value corresponds with the duration of a single latency trial test expressed in seconds. For each frame size, the RFC 2544 algorithm measures the latency several times. This parameter configures the duration of each latency measurement
Iterations	This parameter corresponds with the number of times the latency is measured for every frame size.

7.1.3. Frame Loss

The aim of this test is to determine the frame loss ratio through the entire range of input data rates and frame sizes. The procedure is the following:

1. Send a certain number of frames at a specific rate through the DUT / SUT, counting the frames transmitted and received and computing the frame loss ratio. The first trial should be run for the frame rate that is 100% of the maximum rate supported by the interface.
2. Repeat the procedure for the speed that corresponds to the next test bit rate.
3. Continue this sequence (reducing the bit rate in every step) until there are two consecutive trials where no frames are lost.

Table 7.3: RFC 2544 Frame Loss Test Settings

Setting	Description
Enable	Enables or disables the RFC 2544 frame loss test. The frame loss test will not be executed if is not previously enabled by this control.

Table 7.3: RFC 2544 Frame Loss Test Settings

Setting	Description
Resolution (%)	Measures the decrease in terms of bit rate the tester uses to measure frame loss expressed as a percentage of the nominal transmission capacity. For example if throughput result is 100% for one specific frame size and this field is set to 10%, then frame loss will be evaluated for 100%, 90%, 80%... of the transmission medium capacity. If frame loss is configured to 1%, the bit rates used in consecutive trials will be 100%, 99%, 98%... Setting an smaller *Resolution (%)* increases frame loss test accuracy but the tame it takes to finish the test is also increased.
Trial duration (s.)	This value corresponds with the duration of a single frame loss trial test expressed in seconds. For each configured frame size and load, the equipment sends and analyses traffic for a time period specified by this field. Longer trial duration periods tend to give more accurate frame loss results but they make measurements longer.

7.1.4. Back-to-Back Frames

A back-to-back frames test determines the *node buffer capacity* by sending bursts of traffic at the highest theoretical rate, and then measuring the longest burst where no packets are dropped. The test procedure is as follows:

1. Send a burst of frames with minimum inter-frame gaps to the DUT / SUT, and count the number of frames forwarded.
2. If the count of transmitted frames is equal to the number of frames forwarded, increase the length of the burst and re-run the test. If the number of forwarded frames is less than the number transmitted, reduce the length of the burst and re-run the test.
3. Continue until the back-to-back frames results has been computed with acceptable accuracy.

The back-to-back value is the number of frames in the longest burst that the DUT / SUT can handle without losing any frames. It is recommended to run the test for at least

2 seconds, and it should be repeated at least 50 times with the average of the recorded values being reported.

Table 7.4: RFC 2544 Back-to-back Test Settings

Setting	Description
Enable	Enables or disables the RFC 2544 back-to-back frames test. The back-to-back frames test will not be executed if is not previously enabled by this control.
Maximum burst length (fr.)	This is the maximum frame burst to be used in an back-to-back frames test trial. The DUT / SUT should drop some frames when it receives a burst with the length configured in this field. If the actual back-to-back frames result is larger than the *Maximum burst length (fr.)*, the equipment will be unable to find it.
Resolution (fr.)	It is the minimum back-to-back frames result the equipment distinguishes as different results. When a test cycle finishes, the measurement algorithm compares the burst length with the result of the previous cycle. If the difference between them is smaller than the resolution, then the algorithm terminates. The final back-to-back result is the one computed in the last cycle.
Iterations	Configures the number of bursts sent per test cycle. Increasing *Iterations* makes the back-to-back test result more reliable but it also increases the overall testing time.

7.1.5. System Recovery

This test determines the node speed at which the DUT / SUT recovers from an overload condition. The procedure is as follows:

1. Measure the throughput for the DUT / SUT at each of the listed frame sizes.
2. Send a stream of frames at a rate that is 110% of the recorded throughput rate or the maximum rate for the media, whichever is lower, for at least 60 seconds.
3. At Timestamp A, reduce the frame rate to 50% of the above rate and record the time of the last frame lost (Timestamp B). The system recovery time is calculated by subtracting Timestamp B from Timestamp A. The test must be repeated a number of times, and the average of the recorded values is reported.

The system recovery results should be reported as a table, with a row for each of the tested frame sizes. There should be columns for the frame size, the frame rate used as

the throughput rate for each type of data stream tested, and for the measured recovery time for each type of data stream tested.

Table 7.5: RFC 2544 System Recovery Test Settings

Setting	Description
Enable	Enables or disables the RFC 2544 system recovery test. The system recovery test will not be executed if is not previously enabled by this control.
Trial duration (s.)	Time, computed in seconds, the DUT / SUT is overloaded by frame transmission above the recorded throughput value.
Iterations	Number of times the recovery time is measured in one test cycle. The final value of the recovery time will be the averaged value computed for each iterations. The more iterations per test cycle the more reliable the test results are. However, increasing the iterations increases the total testing time as well.

7.1.6. Configuration Pass / Fail Thresholds

One of the advantages of the RFC 2544 tests is the ability to supply a clear pass / fail results based on certain thresholds configured before the test start. To configure the pass / fail thresholds follow these steps:

Table 7.6: RFC 2544 Performance Objectives

Setting	Description
Throughput	Sets a performance objective for the RFC 2544 throughput test. The test will be considered to fail if the measured throughput is smaller that the threshold and it will be passed otherwise.
Latency	Sets a performance objective for the RFC 2544 latency test. The test is considered to fail if the result is larger than the threshold and it is considered to pass otherwise.
Frame loss	Configures an performance objective for the RFC 2544 frame loss test. The test is considered to fail if there is at least one result within the frame length and bit rate ranges which produces a frame loss figure larger than the threshold. The test is considered to pass otherwise.

Table 7.6: RFC 2544 Performance Objectives

Setting	Description
Back-to-back frames	Sets a performance objective for the RFC 2544 back-to-back frames test. The test is considered to fail if the back-to-back frames result is smaller than the configured threshold and it is considered to pass otherwise.
System recovery	Sets the performance objective for the RFC 2544 system recovery time test. The test fails if the system recovery time result is longer than the threshold and passes otherwise.

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Go to *Performance objectives*,
3. Select *RFC 2544*.
   Note: You must enable the RFC 2544 test before you are allowed to enter in this menu.
4. Enter the thresholds for *Minimun throughput (%)*, *Maximum latency*, *Maximum frame loss (%)*, *Minimum frame burst (fr)*, *Maximum recovery time*.
   Note: Each specific test has to be individually enabled before you can set the performance objective for it.

7.1.7. Getting Test Results

The RFC 2544 provides a description of the way test results must be presented to the user. Normally, these results are displayed as tables. Each row represents the test frame length or bit rate and columns are usually test results like throughput, latency or others.

The RFC 2544 results can be evaluated as pass or fail. Failed results are represented with red colour. One test is considered to fail it at least one test result fails. The whole RFC 2544 fails if one particular test result fails. The global Pass / Fail can be checked at any moment, even before the end of the test from the *Summary* panel (SUM key).

Once the test has been configured, it can be started at any time with the help of the RUN button. You can wait to the test to finish but it is also possible to check partial test results at any moment. To do that follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
   Note: If you have configured *Test method* to *One-way (A > B),* the RFC 2544 results are available in *Port B*. On the other hand, if you have configured *Two-way (A > A)*, the RFC 2544 results are available in *Port A*.
3. Select *RFC 2544*.

4. Check the *Status* field to know the pass / fail test result. If the test has not yet finished, this field will display a *In progress* message. If the equipment detects an error during text execution it will display an error message.
5. Check the *Current stage* and *In progress* fields to know what is doing the tester in the current moment and which is the current test progress.

Figure 7.1: RFC 2544 test results: (a) Throughput results table, (b) Latency results table.

6. Go to the *Throughput test*, *Latency test*, *Frame lost test*, *Back-to-back test* and *System recovery test* and check the detail led results for each of them.

(a)

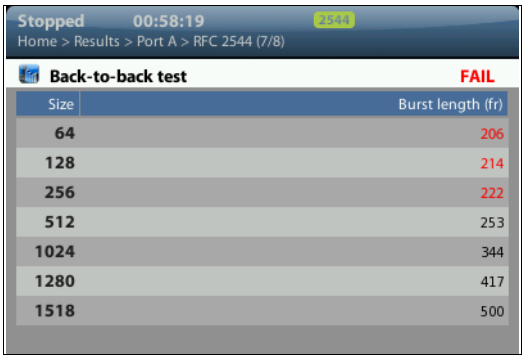

(b)

(c)

Figure 7.2: RFC 2544 test results: (a) Frame loss result table, (b) Back-to-back frames result table, (c) System recovery time results table.

Note: Test results depend on the frame length and some of them also on the throughput. For this reason, RFC 2544 result tables depend on both magnitudes: *Size* and *Throughput (%)*. *Throughput (%)* represents the bit rate used for testing latency and frame loss. It can be the previously measured maximum rate attainable bit rate expressed as a percentage of the nominal medium rate (10 Mb/s, 100 Mb/s, 1000 Mb/s) or a different value depending on the particular test settings.

Note: The system recovery time tests sends traffic above the throughput value in order to overload the DUT / SUT, the test traffic in this case is indicated in the results table by means the *Test rate (%)* column.

Table 7.7: RFC 2544 Test Results

Setting	Description
Theoret. max rate (fr/s)	Theoretical maximum rate attainable in the transmission medium expressed in frames per second. In an Ethernet link, this rate is calculated as nominal transmission capacity in bits per second (10 Mb/s, 100 Mb/s, 1000 Mb/s) divided by the frame bits (including the 64-bit preamble and the 96-bit inter-frame gap). The theoretical maximum rate result is available in the *Throughput test* result table.
Measured max rate (fr/s)	Measured maximum rate attainable in the transmission medium expressed in frames per second. This result is always smaller or equal than the theoretical maximum rate result. The measured maximum rate result is available in the *Throughput test* result table
Delay (µs)	Latency result measured in microseconds. The delay result has to be understood as the one-way delay if the test is configured as a *One-way (A > B)* test or as the round-trip delay if the test is configured as *Two-way (A > A)*. In many setups the round-trip delay is roughly twice that the one-way delay but in asymmetric setups this is not true anymore.
Frame loss (%)	Frame loss result expressed as a percentage of the total amount of frames transmitted in the trial for an specific frame length and throughput.
Burst length (fr)	Longest burst accepted by the DUT / SUT which produces no packet loss.
Recovery time (µs)	Time invested to recover from an overload condition expressed in microseconds.

7.1.8. Generation of RFC 2544 Result Reports

An essential part of standardized network testing is report generation. Reports summarize test results in a document that can be edited or shared with the customer. RFC 2544 reports include a text header with basic information about the test like start time, duration and configuration. Test results are presented in tables. There is one table for each test: throughput, latency, frame loss, back-to-back frames and system recovery time. All results include a Pass / Fail indication. The report also includes a global Pass / Fail indication which summarizes all test results: It follows an example of a typical RFC 2544 report as is generated by Ether.Genius / Ether.Sync / Ether.Giga:

```
RFC 2544 Test Report

Report name             2012-11-22-185949
Customer
Department
Company
Location
Operator
Start Time              Thu Nov 22 18:59:49 2012
Elapsed Time            11:16:44
Test Unit               Ether.Giga Gigabit Ethernet Tester
Serial number           MEM0009P
Software version        0.9.3


Global results

Status                              FAIL
Completed                           100 %


Test Unit Configuration

Mode                                IP endpoint
Test method                         Two-way (A > A)

                                                  Port A              Port B
Port mode                                         TX/RX            Loopback
Connector                                     Electrical          Electrical
Encapsulation                                       None
Source MAC address                    00:DB:1E:00:01:10
Destination MAC address               00:DB:1E:00:01:11
Address range size                                  ----
C-VID                                               ----
C-VLAN priority                                     ----
S-VID                                               ----
S-VLAN priority                                     ----
DEI                                                 ----

Source IPv4 address                          172.26.3.23
Destination IPv4 address                     172.26.4.24
Destination host name                               ----
Address range size                                  ----
DSCP                                                   0

Performance objectives
```

Minimum throughput (%)	50.000
Maximum latency	10.000 ms
Maximum frame loss (%)	1.000
Minimum frame burst (fr)	1000
Maximum recovery time	10.000 ms

Throughput test

Frame sizes	Theor.max (fr/s)	Max.rate (fr/s)	Max.rate (%)	Status
64	148,809	40,399	27.140	FAIL
128	84,459	39,755	47.070	FAIL
256	45,289	39,009	86.130	PASS
512	23,496	23,496	100.000	PASS
1024	11,973	11,973	100.000	PASS
1280	9,615	9,615	100.000	PASS
1518	8,127	8,127	100.000	PASS

Latency test

Frame sizes	Throughput (%)	Delay (us)	Status
64	27.140	778.03	PASS
128	47.070	84.47	PASS
256	86.130	228.79	PASS
512	100.000	172.88	PASS
1024	100.000	140.27	PASS
1280	100.000	162.50	PASS
1518	100.000	185.06	PASS

Frame loss test

Throughput (%)	64	128	256	512	1024	1280	1518
100.00	73.126	52.885	13.869	0.000	0.000	0.000	0.000
90.00	70.086	47.603	4.368	0.000	0.000	0.000	0.000
80.00	66.364	40.995	0.000	---	---	---	---
70.00	61.527	32.571	0.000	---	---	---	---
60.00	55.078	21.274	---	---	---	---	---
50.00	46.084	5.506	---	---	---	---	---
40.00	32.623	0.000	---	---	---	---	---
30.00	10.208	0.000	---	---	---	---	---
20.00	0.000	---	---	---	---	---	---
10.00	0.000	---	---	---	---	---	---
	FAIL	FAIL	FAIL	PASS	PASS	PASS	PASS

Back-to-back test

Frame sizes	Burst length (fr)	Status
64	46,875	PASS
128	46,875	PASS
256	46,875	PASS
512	3,000,000	PASS
1024	3,000,000	PASS
1280	3,000,000	PASS
1518	3,000,000	PASS

System recovery test

Frame sizes	Throughput (%)	Recovery time	Status

```
   64          27.14          620.400 us                    PASS
  128          47.07          906.000 us                    PASS
  256          86.13          335.000 us                    PASS
  512         100.00            0.000 us                    PASS
 1024         100.00            0.000 us                    PASS
 1280         100.00            0.000 us                    PASS
 1518         100.00            0.000 us                    PASS

                                (c) 2012 ALBEDO Telecom
```

7.2. Performance Assessment with the eSAM Test

Ethernet service activation though eSAM defined in ITU-T 1564 has arisen as an alternative to RFC 2544 verification. Unlike the RFC 2544, eSAM is designed for Ethernet service activation from the beginning. The advantages of eSAM in front of RFC 2544 are summarized in the following points:

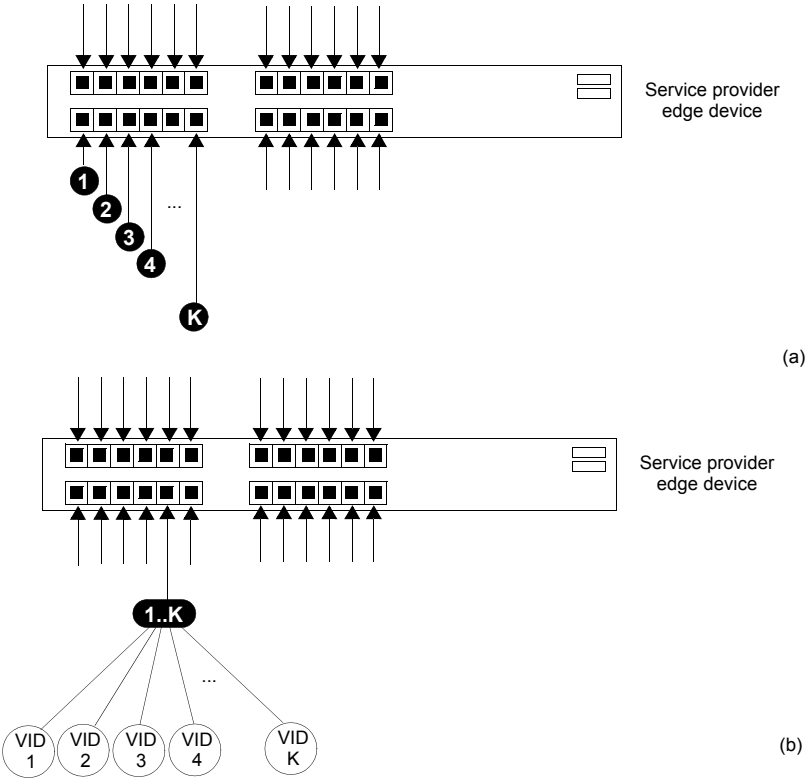

Figure 7.3: (a) Port based Ethernet service, one service per port. (b) Ethernet service multiplexing based on service delimiting markers like the VLAN tag.

1. *Faster execution*: An standard eSAM test is made up of a short configuration test and a longer performance test. If the configuration test fails there is no need to execute the long performance test. The result is that network administrators have time to correct any configuration issue before having to wait for the complete test execution.
2. *FDV results*. Frame Delay Variation (FDV) is a key metric to evaluate network performance. FDV is very sensitive to congestion and other degradations that affect end-to-end network performance and it is therefore an essential parameter to measure.

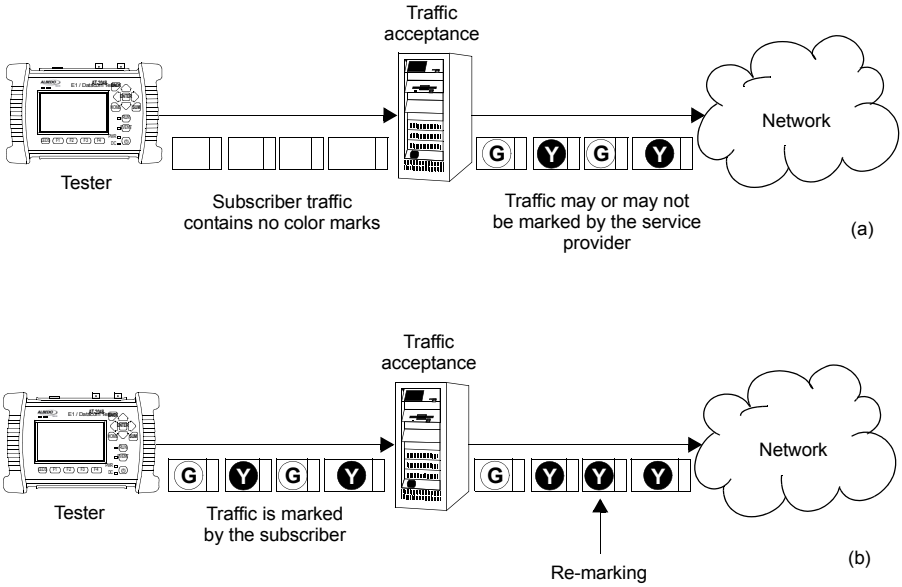

Figure 7.4: (a) Non-color-aware service. It the service provider is wiling to implement different transmission priorities for this service, marking will have to be carried out by the access network by means some traffic classification algorithm. (b) In a color-aware service subscribers mark their own traffic from the beginning. The service provider may remark some traffic depending on the traffic acceptance algorithm.

3. *Compatible with multiservice environments*: Modern Ethernet services may be port based but service multiplexing in the same port by means some service delimiting tag is also very popular. The eSAM test has been designed to operate in environments using service multiplexing. In this case, all services are simultaneously tested and independent results are given for each of them.

4. *Supports color-aware traffic*: The eSAM test is compatible with color markers used by some service providers to enable the different performance levels in their applications. Color markers classify the network traffic in three sets: *green* traffic is transmitted within the delay and frame loss ratio limits guaranteed by the SLA agreement, *yellow* traffic is transmitted but the SLA agreement performance limits do not apply for it, and finally *red* traffic is discarded and not transmitted (red traffic is therefore never seen in the network). Common color markers used in practical applications are the DSCP (Layer 3) and the VLAN priority bits (Layer 2). The latter requires the subscriber frames to be encapsulated in VLAN tagged frames.

7.2.1. Bandwidth Profiles for Ethernet Services

To see how eSAM works it is essential to understand how Ethernet services are defined. Network operators have at their disposal the tools that enable them to define their services with great flexibility. The information rate associated with Ethernet service is not limited to the nominal speed of the access network interface. For example, certain operator way want to define a 2 Mbit/s service over an optical Gigabit Ethernet interface.

Performance in terms of delay, packet loss and other metrics is applied to traffic flows defined by their generation statistics or bandwidth profile. The mechanism used by service providers to make sure the ingress traffic has the correct bandwidth profile is admission control. Once the Ethernet access has been set up, the service provider performs admission control over the customer traffic at the user-network interface. Admission control for Ethernet services uses bandwidth profiles based on four parameters initially defined by the Metro Ethernet Forum (MEF):

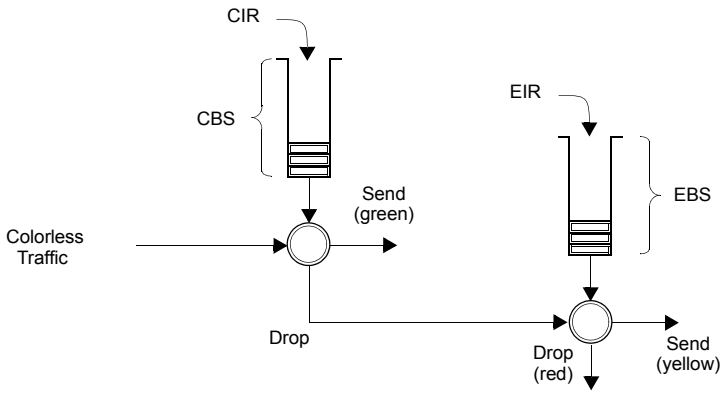

Figure 7.5: Two-rate Three color marker (trTCM). policing algorithm.

- *Committed Information Rate (CIR)*: Rate up to which service frames are delivered as per the service performance objectives.

- *Committed Burst Size (CBS)*: Maximum number of bytes up to which service frames may be sent as per the service performance objectives without considering the CIR.
- *Excess Information Rate (EIR)*: Rate up to which service frames are still delivered but they are not subject to any performance objective.
- *Excess Burst Size (EBS)*: The number of bytes up to which service frames are sent (without performance objectives), even if they are out of the EIR threshold.

The MEF specifies the RFC 2698 *Two-rate Three-Color Marker* (trTCM) as the admission control mechanism for carrier Ethernet services. The trTCM is obtained by chaining two simple token bucket policers. Tokens fill the main bucket until they reach the capacity given by the CBS parameter, at a rate given by the CIR parameter. The secondary bucket is filled with tokens with the EIR rate until they reach the capacity given by the EBS parameter.

The traffic that passes through the first bucket (green traffic) is delivered with the QoS agreed with the service provider, but any traffic that passes through the secondary bucket (yellow traffic) is re-classified and delivered as best-effort traffic, or it is given a low priority. Non-conforming traffic (red traffic) is dropped.

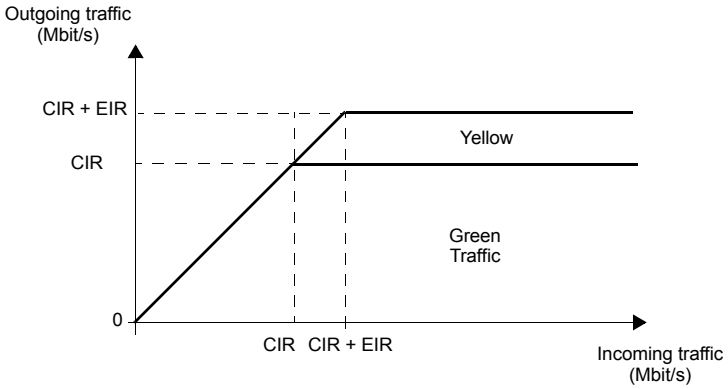

Figure 7.6: The amount of traffic that crosses an admission control filter. Graphics represent steady states, traffic is usually allowed to be greater than the CIR and EIR for short periods of time. Traffic delivery is guaranteed if the rate is smaller than the CIR. Excess traffic (EIR traffic) is delivered as well, but it is marked as low priority and usually discarded first if congestion occurs.

Note than the best effort classical service can be obtained simply by setting the CIR parameter to zero. Moreover, service providers may allow their subscribers to add their color marks to the traffic they generate before the traffic admission algorithm is applied. This kind of pre-marking transfers more control on the application performance to the

end user. However, admission control is necessary even in this case and frame remarking is done on non-conforming traffic anyway.

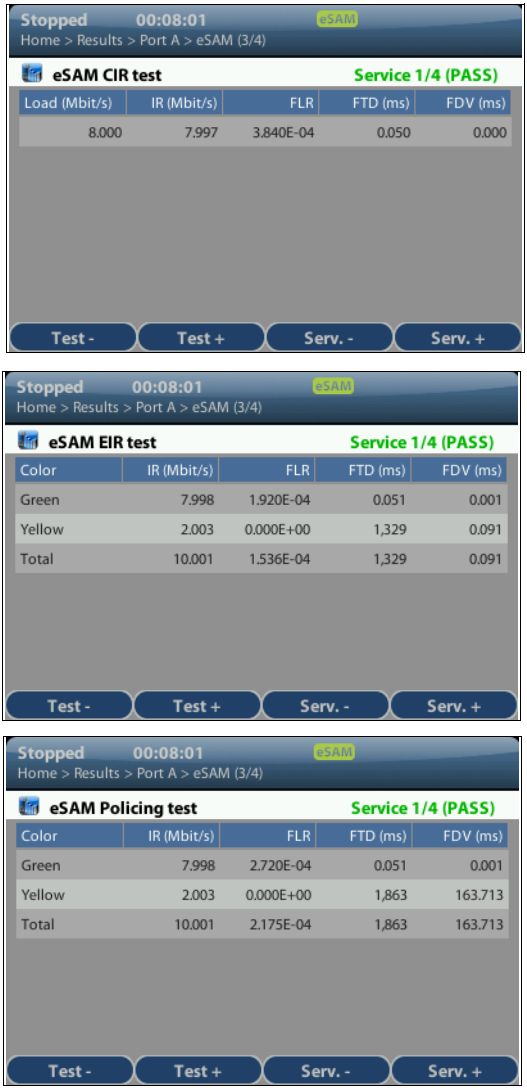

Figure 7.7: CIR, EIR and policing test results as presented by Ether.Genius / Ether.Sync / Ether.Giga.

The basic purpose of eSAM is to check that green and yellow frames are transported with the required performance in terms of Frame Total Delay (FTD), Frame Delay

Variation (FDV), Frame Loss Ratio (FLR) and availability. The existence of a policing algorithm like the trTCM means that testing the ability of the network to discard non-conforming traffic is another important requirement for eSAM. Transmission of green traffic is verified by the CIR test and the performance test, transmission of the yellow traffic is checked with the EIR test and red (discarded) traffic is measured by the policing test. Finally, the Information Rate (IR) is measured in all CIR, EIR and policing tests for all traffic classes to make sure that the information is preserved by the network when required to do so. More details about these specific tests are provided in the following sections.

7.2.2. Test Configuration

Like almost any test to be carried out by Ether.Genius / Ether.Sync / Ether.Giga, the way the equipment is configured depends on the particular network and service to be tested. Before running the test, there are several questions the user must answer:

- What kind of network is going to be tested? Configuration is different for IP networks and Ethernet networks. Note than a network may carry IP over Ethernet frames but it may still be more interesting to run an Ethernet test than an IP test.
- How is the test equipment going to be connected to the network? Configuration is not the same if there is a traffic reflector used to loop frames back to the analyser (two-way test) or if Port B is going to be used for analysis (one-way test).
- How many services are required to be measured in the same test? Is there any color marker used to classify the traffic? Which are the CIR / EIR values for each service? Is there any policing mechanism for the Ethernet services configured in the network?
- What is the required performance for each service in terms of FTD, FDV, FLR and availability?
- In IP tests, what are the correct IP profiles to be used? IP addresses, network masks, gateways and DNS servers (if used) must be known before the test can be configured. Usually DHCP protocol makes easier configuration of IP profiles but DHCP may not be available in some networks.

Table 7.8: eSAM Configuration

Setting	Description
Color mode	Set this field to *On* if the traffic contains color labels. Color labels enable the network to supply differentiated services to selected traffic. These services consist in low delay paths, high priority frame scheduling or low packet loss probability. Set the *Color mode* to *Off* if you want to leave the service provider network to decide the priority of user traffic based on the result of a traffic acceptance algorithm (policer, shaper).

Table 7.8: eSAM Configuration

Setting	Description
Color method	The color method configures which field in the test traffic implements the color mechanism. There are the following potential choices: • *DSCP*: Differentiated Services Code Point. It is 6-bit class of service label that accepts values between 0 and 63. DSCP makes sense as a color marker in routed networks and for this reason it is not available in *Ethernet endpoint mode*. • *C-VLAN priority*: 3-bit class of service field carried by the VLAN tag in IEEE 802.1Q frames or in the C-VLAN tag in IEEE 802.1ad and Q-in-Q frames. Using this field requires frames containing at least one VLAN tag. Color markers based on the C-VLAN priority field make sense only in switched networks for this reason it is not available in *IP Endpoint*. • *C-VID*: VLAN identifier assigned to tagged frames (IEEE 802.1Q) or C-VLAN identifier for double tagged frames (IEEE 802.1ad, Q-in-Q). The VID is commonly used in carrier Ethernet networks as a service discriminator. The Ether.Genius / Ether.Sync / Ether.Giga testers account for the possibility to use this field as a colour marker as well. It requires the tester to be configured in *Ethernet Endpoint* mode and an encapsulation with at least one VLAN tag. • *S-VLAN priority*: 3-bit class of service field carried by the S-VLAN tag in IEEE 802.1ad and Q-in-Q frames. Color markers based on the C-VLAN priority field make sense only in switched networks for this reason it is not available in *IP Endpoint* mode . • *S-VID*: VLAN identifier assigned S-VLAN tag in double tagged frames (IEEE 802.1ad, Q-in-Q). The VID is commonly used in carrier Ethernet networks as a service discriminator. The Ether.Genius / Ether.Sync / Ether.Giga testers account for the possibility to use this field as a colour marker as well. It requires the tester to be configured in *Ethernet Endpoint* mode and an IEEE 802.1ad / Q-in-Q encapsulation
Number of services	Defines the number of services to be simultaneously tested. The maximum number of services to be tested is eight (non-color-aware mode) or four (color-aware mode).

Table 7.8: eSAM Configuration

Setting	Description
Number of steps	Number of different bit rates to be tested in the configuration CIR test. The test bit rates are equally distributed between 0 and CIR bit rates.
Step duration (s)	Duration of each iteration of the CIR test. If Number of steps is configured to 1, it is the duration of the CIR test. Any value between 1 and 60 seconds is accepted.
Policing test	Enable to execute the policing test as a part of the eSAM test suite or disable if you don't want to run the policing test. You may want to disable the policing test if you know that the network under test is not using any traffic admission algorithm and you don't want the configuration test to fail for this reason.
Performance test duration	Duration of the eSAM performance test. The performance test is executed once the configuration test has finished. For the performance test you can choose between one of the three duration presets (*15 min.*, *2 hours*, *24 hours*) or you can set your own test duration with a resolution of one second by setting this field to *User duration*. It is possible to use this field to disable the performance test. This is interesting if you want to check the network configuration but you don't need to know the performance.
User duration	Use this field to configure the eSAM performance test duration if you have set *Performance test duration* to *User duration*.
Service configuration	Displays a panel that enables the user to enter the CIR and EIR for each service to be tested. The service configuration panel also displays information about the Ether.Genius / Ether.Sync / Ether.Giga traffic flows assigned to each eSAM service. Two flows are required to test one color-aware service and one flow is necessary for each non-color-aware service.

Once all the details about test equipment connection and network / service configuration have been clarified is time to configure the test. To do that follow these steps:

1. Make sure that your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).

2. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
3. Choose between *Ethernet endpoint* or *IP endpoint* with the *Mode* setting.
4. Depending on your test setup (See section 2.2.3), configure *Test method* to *One-way (A > B)* or *Two-way (A > A)*.
5. Configure *Performance test* to *eSAM*.
6. Enter in the *eSAM* menu to configure the global ITU-T Y.1564 settings.
7. Configure the *Color mode* to *On* if the service provider allows pre-marking of Ethernet frames / IP datagrams or set it to *Off* otherwise.
8. If you have configured *Color mode* to *On*, set the *Color method* to *DSCP* (IP Endpoint mode), *C-VLAN priority*, *C-VID*, *S-VLAN priority* or *S-VID* (Ethernet Endpoint mode).
   Note: C-VLAN priority and C-VID color markers require VLAN, Q-in-Q or IEEE 802.1ad encapsulation. S-VLAN priority and S-VID color markers require Q-in-Q or IEEE 802.1ad encapsulation
9. Set the number of services you want to test within the same test with the help of the *Number of services* control.
   Note: The maximum number of services is 4 if *Color mode* is *On* or 8 if *Color mode* has been configured to *Off*.
10. Configure the *CIR* and *EIR* parameters for each service to be tested from the *Service configuration* panel.
11. Enter the number of steps and duration of each step for the CIR test with the help of the *Number of steps* and *Step duration (s)* controls.

Table 7.9: eSAM Performance Metrics

Result	Description
Load (Mbit/s)	Amount of test traffic offered to the network expressed in Mbit/s. This traffic is always larger or equal than the measured Information Rate (IR). Currently this parameter is displayed as a result only for the CIR configuration test. Traffic load for EIR and performance tests can be checked in the service configuration panel for eSAM tests. For policing tests, the load is always higher than the CIR + EIR in an amount than is computed automatically as specified in ITU-T Y.1564.
IR (Mbit/s)	Average Information Rate (IR) computed for the traffic currently being tested (green, yellow, aggregated, etc.) within the configured test period. Only test traffic is taken into account to measure the IR. The definition and application of the IR metric follows ITU-T Y.1564.

Table 7.9: eSAM Performance Metrics

Result	Description
FLR	Ratio of the total amount of lost frames to the total transmitted frames from the beginning of the test. Definition of the FLR parameter follows ITU-T Y.1563.
FTD (ms)	Is the worst case (maximum) Frame Total Delay (FTD) found from the beginning of the eSAM configuration or performance test and expressed in milliseconds. The definition of *FTD (ms)* follows ITU-T Y.1563 and it is computed in the same way that the *FTD (maximum)* metric supplied by the Ether.Genius / Ether.Sync / Ether.Giga *SLA statistics*.
FDV (ms)	Is the worst case (maximum) Frame Delay Variation (FDV) found from the beginning of the eSAM configuration or performance test and expressed in millisecond. The definition of *FDV (ms)* follows RFC 3393 and RFC 1889 and it is computed in the same way that the *FDV (maximum)* metric supplied by the Ether.Genius / Ether.Sync / Ether.Giga *SLA statistics*.
Avail (%)	The Avail (%) constitutes an availability performance figure that informs about the percentage of time that the DUT / SUT has been not available for transmit / receive data during the eSAM performance test. This performance metric is computed as *100% - PEU(%)*. Where the PEU(%) is the Percent Ethernet service Unavailability defined in ITU-T Y.1563 and supplied by the Ether.Genius / Ether.Sync / Ether.Giga *SLA statistics*. This metric is computed for eSAM performance tests only. Configuration tests are considered too short to make any accurate availability result significant enough.

12. Enable the Policing test with *Policing test* if your network is using a traffic admission mechanism that limits the amount of accepted ingress traffic
13. Configure the duration of the eSAM performance test with the help of the *Performance test duration* and *User duration* controls.
14. Leave the *eSAM* configuration panel and from the *Test* menu select Performance objectives.
15. Select *eSAM* and enter the performance objectives in terms of the *FLR*, *FTD*, *FDV* and *Avail.* fields for each service you want to test.

16. Configure the *Frame layer* (See section 4.1.2) in Port A for all your services. Parameters to be configured are source and destination MAC addresses, VLANs, frame size, etc.
    Note: To know the correspondence between the Ether.Genius / Ether.Sync / Ether.Giga flows and the eSAM services, check the *Service configuration* panel you have used to configure the CIR and EIR values for the test.
    Note: Test traffic corresponding to different services should not have exact configurations. Otherwise, the analyser (and the network) will fail to classify the traffic. A service delimiting field could be used for this purpose. It is common to use the VID but anything that makes traffic from different services different in some way is accepted by the tester.
    Note: If you are configuring colored traffic, the color makers (C-VLAN priority, C-VID, S-VLAN priority, S-VID) must have a different value for green and yellow traffic corresponding to the same service. The test will fail to start if this requirement is not met.
17. If you are working in *IP endpoint* mode, configure the *Local profile* (See section 2.3) and the *Network layer* (See section 4.4.4) for all your services. Parameters to be configured are IP addresses, DSCPs, etc.
    Note: Previous notes about flows / services, service delimiting tags and color markers are valid in *IP endpoint* as well but in this case the color marker to be used is the DSCP rather than VLAN priorities or VIDs.
18. Run the eSAM test with the help of the *RUN* fixed button.

7.2.3. CIR Configuration Test

The CIR Configuration test is the most basic eSAM test. Its purpose is to check the ability of the network to deliver frames at CIR rate within acceptable performance limits. This test consist in loading the service with the maximum bit rate it supports without any degradation that is the CIR by definition. Optionally, the user is allowed to configure other test rates smaller than the CIR. For colored services, the only traffic color to be generated in the CIR test is green.

CIR test results are made up of the IR, FTD, FDV and FLR for all the test loads. The test is considered to pass if the computed performance metrics (FTD, FDV and FLR) are better than the corresponding performance thresholds. Test results, including the final Pass / Fail are essentially the same for color aware and non-color aware CIR tests.

It is possible to execute several (up to four color-aware or eight non-colour-aware) CIR tests. In this case, CIR tests are executed sequentially for each configured service. To check the CIR test results follow this procedure:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
   Note: If you have configured *Test method* to *One-way (A > B),* the eSAM results

are available in *Port B*. On the other hand, if you have configured *Two-way (A > A)*, the RFC 2544 results are available in *Port A*.

3. Select *eSAM*.
4. Check the *Status* field to know the pass / fail test result. If the test has not yet finished, this field will display a *In progress* message. If the equipment detects an error during text execution it will display an error message.
5. Check the *Test remaining time* to get an estimation about how much time is left to finish the current test.
6. Go to *eSAM Configuration test*.
   A table with eSAM configuration test results is displayed.
7. Use the *Test-* (*F1*) and *Test+* (*F2*) contextual keys to select the eSAM CIR test results. Use the *Serv-* (*F3*) and *Serv+* (*F4*) contextual keys to select the table corresponding to the service you want to check.

7.2.4. EIR Configuration Test

The purpose of the EIR Configuration test is to measure network performance when it is loaded with an information rate that matches the CIR + EIR. Unlike it happens with the CIR configuration test, the way the results are computed is very different for color-aware and non-color-aware services. The reason is that there is no performance limit in non-colour-aware services when the transmission rate is above the CIR but for color-aware services it is still possible to guarantee the quality of service of green frames. The EIR test is executed and evaluated as follows:

- *Non-color-aware services*: The network is loaded with a CIR + EIR bit rate and the IR, FTD, FDV and FLR are measured. The test is considered to pass if CIR * (1 - FLR) < IR < CIR + EIR.
- *Color-aware services*: The network is loaded with CIR green traffic and EIR yellow traffic. The IR, FTD, FDV and FLR is measured for both traffic classes. The test is passed if the FTD, FDV and FLR for green traffic are within acceptable limits.

It must be noticed that performance metrics are always measured but they are relevant for the pass / fail results only for guaranteed (green) traffic.

It is possible to execute several (up to four color-aware or eight non-colour-aware) EIR tests. In this case, EIR tests are executed sequentially for each configured service. Users may not want to run the EIR test. To do that, they have simply to configure the EIR to 0 for services where the test is not going to be executed.

The procedure to follow to display the EIR test results is similar than for the CIR results but in this case it is necessary to use he *Test-* (*F1*) and *Test+* (*F2*) contextual keys to select the eSAM EIR test results rather than the CIR ones

7.2.5. Policing Configuration Test

The Policing test is useful to make sure that the network drops non-conforming (red) traffic. The policing test loads the network with a policing rate computed automatically

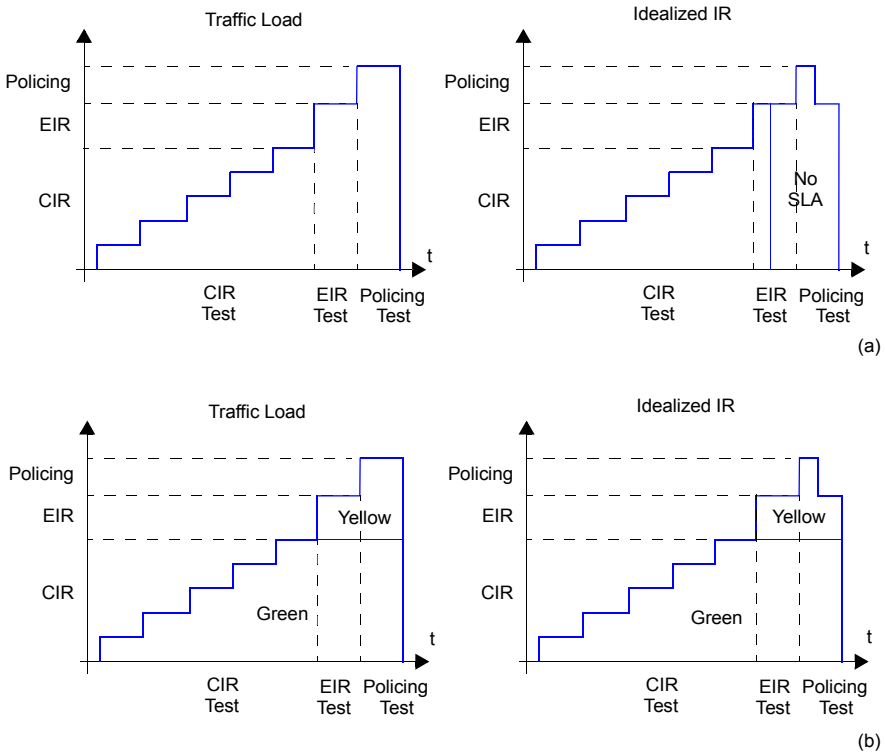

Figure 7.8: Typical CIR, EIR and Policing tests in colour-aware and non colour aware services: (a) No colours are defined in the interface. Some traffic may be degraded if CIR < IR < EIR because quality of service is not guaranteed for excess traffic. If IR > EIR some traffic will be lost due to the action of the policing filter. (b) Coloured interface: Yellow traffic does not have quality of service guarantees but SLA is meet for Green traffic as long as the IR remains smaller than the CIR.

as specified in ITU-T Y.1564. The policing test is always higher than the sum of the CIR and the EIR. Again, test execution and result presentation is different in color-aware and non-color aware services. Details are as follows:

- *Non colour-aware services*: The network is loaded with a CIR + EIR + policing bit rate and the IR, FTD, FDV and FLR are measured. The test is considered to pass if CIR * (1 - FLR) < IR < CIR + EIR + 1%. The extra 1% is used to account for the burstability of policing filters due to an EBS parameter different to zero.
- *Color-aware services*: The network is loaded with CIR green traffic and EIR + policing yellow traffic. The IR, FTD, FDV and FLR is measured for both traffic classes. The test is passed if the FTD, FDV and FLR for green traffic are within

acceptable limits and if the aggregated IR (green + yellow) meets the following double inequality: CIR * (1 - FLR) < IR < CIR + EIR + 1%.

Like it happens with the CIR and EIR tests, it is possible to execute several (up to four color-aware or eight non-colour-aware) policing tests. In this case, policing tests are executed sequentially for each configured service.

Users may choose to disable the policing test if the network is not using any admission control mechanism based on policing filters. The policing test fails if it is executed in networks not supporting policing.

The procedure to follow to display the policing test results is similar than for the CIR results but in this case it is necessary to use he *Test-* (*F1*) and *Test+* (*F2*) contextual keys to select the eSAM policing test results rather than the CIR ones

7.2.6. Performance Test

The eSAM performance test is executed only if the configuration CIR, EIR and policing tests are passed. While configuration tests take a few minutes to finish, a performance test may take hours or even days depending on the test requirements.

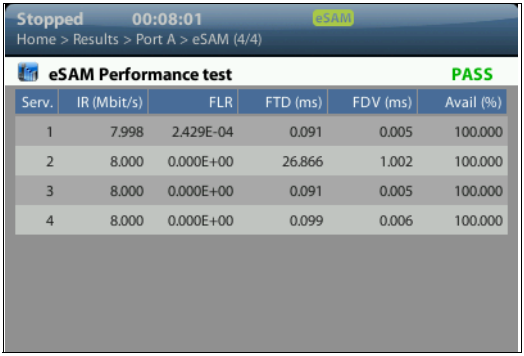

Serv.	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Avail (%)
1	7.998	2.429E-04	0.091	0.005	100.000
2	8.000	0.000E+00	26.866	1.002	100.000
3	8.000	0.000E+00	0.091	0.005	100.000
4	8.000	0.000E+00	0.099	0.006	100.000

Figure 7.9: Performance results as presented by Ether.Genius / Ether.Sync / Ether.Giga.

In some ways, the performance test is similar to the CIR test because in the network is loaded with the CIR bit rates of all configured service but in this case all services are tested simultaneously and therefore the total load is the sum of all CIR rates for all Ethernet services under test. This is important because a network may be able to support all services if they are not all them loaded at the same time but it may fail to support all services operating at maximum speed. The second difference between CIR and performance tests is that performance test uses to be long enough to compute a significant availability figure.

The performance test computes the IR, FTD, FDV, FLR and availability for all the services being tested. The test is passed if all performance metrics are found to be within acceptable limits.

The procedure to follow to display the performance test results is similar than for the CIR results but in this case it is necessary to go the *eSAM Performance test* result panel rather than to the *eSAM Configuration test* panel.

7.2.7. Generation of eSAM Result Reports

Report generation for eSAM works in a similar way that for RFC 2544. There's no need to make any special configuration in order to get the eSAM test report. It is only necessary to enable report generation (See section 11.1).

Report organization is similar than in other tests, The report contains a header, containing details about which equipment has been used to run the test, when the tester has been executed, and other information, a global PASS / FAIL indication, followed by test-specific PASS / FAIL indications, a summary of the test setup and finally detailed descriptions of the eSAM configuration and performance tests.

```
eSAM test report

Report name              2013-04-11-173048
Custom
Department
Company
Location
Operator
Start Time               Thu Apr 11 17:30:48 2013
Elapsed Time             00:03:08
Test Unit                Ether.Giga Gigabit Ethernet Tester
Serial number            MEM0009P
Software version         1.5.3


Global results (Port A)

                                    Service         Status

Global status                           ---           PASS

eSAM Configuration test                   1           PASS
eSAM Configuration test                   2           PASS
eSAM Configuration test                   3           PASS
eSAM Configuration test                   4           PASS
eSAM Performance test                   ---           PASS


Test Unit Configuration

Mode                                      IP endpoint
Test method                               Two-way (A > A)
```

	Port A	Port B
Port mode	TX/RX	Loopback
Connector	Electrical	Optical

eSAM test configuration

Color mode	On
Color method	DSCP
Number of services	4
Number of steps	1
Step duration (s)	3
Policing test	Disabled
Performance test duration	User duration
User duration	00:01:00

Service configuration

Service	CIR	EIR
1	2.000 Mb/s	2.000 Mb/s
2	2.000 Mb/s	2.000 Mb/s
3	2.500 Mb/s	2.500 Mb/s
4	2.500 Mb/s	2.500 Mb/s

eSAM objectives

Service	FLR	FTD	FDV	Avail.
1	1.000E-03	100.000 ms	50.000 ms	99.000 %
2	1.000E-03	100.000 ms	50.000 ms	99.000 %
3	1.000E-03	100.000 ms	50.000 ms	99.000 %
4	1.000E-03	100.000 ms	50.000 ms	99.000 %

eSAM CIR test

Service 1

Load (Mbit/s)	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
2.000	2.000	0.000E+00	0.145	0.014	PASS

Service 2

Load (Mbit/s)	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
2.000	2.000	0.000E+00	0.158	0.015	PASS

Service 3

Load (Mbit/s)	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
2.500	2.500	0.000E+00	0.149	0.022	PASS

Service 4

Load (Mbit/s)	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
2.500	2.500	0.000E+00	0.148	0.019	PASS

eSAM EIR test

Service 1

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	2.000	0.000E+00	0.395	0.115	PASS
Yellow	2.000	0.000E+00	0.388	0.116	PASS
Total	4.000	0.000E+00	0.395	0.116	PASS

Service 2

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	2.000	0.000E+00	0.141	0.012	PASS
Yellow	2.000	0.000E+00	0.144	0.012	PASS
Total	4.000	0.000E+00	0.144	0.012	PASS

Service 3

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	2.500	0.000E+00	0.498	0.184	PASS
Yellow	2.500	0.000E+00	0.495	0.181	PASS
Total	5.000	0.000E+00	0.498	0.184	PASS

Service 4

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	2.500	0.000E+00	0.196	0.020	PASS
Yellow	2.500	0.000E+00	0.219	0.021	PASS
Total	5.000	0.000E+00	0.219	0.021	PASS

eSAM Policing test

Service 1

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	----	---	----	----	----
Yellow	----	---	----	----	----
Total	----	---	----	----	----

Service 2

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	----	---	----	----	----
Yellow	----	---	----	----	----
Total	----	---	----	----	----

Service 3

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	----	---	----	----	----
Yellow	----	---	----	----	----
Total	----	---	----	----	----

Service 4

Color	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Status
Green	----	---	----	----	----
Yellow	----	---	----	----	----
Total	----	---	----	----	----

eSAM Performance test

Service	IR (Mbit/s)	FLR	FTD (ms)	FDV (ms)	Avail.	Status
1	2.000	0.000E+00	0.501	0.071	100.000 %	PASS
2	2.000	0.000E+00	0.540	0.061	100.000 %	PASS
3	2.500	0.000E+00	0.484	0.063	100.000 %	PASS
4	2.500	0.000E+00	0.532	0.075	100.000 %	PASS

Chapter 8
Ping and Traceroute Tools

Ping and Traceroute are two basic IP network verification tools. Both Ping and Traceroute can be considered as an integrating part of the Operations, Administration and Maintenance (OAM) suite for the IP protocol family. Due to the high availability of this tools, Ping and Traceroute can be used for testing in almost any network.

8.1. Ping

Ping checks “distance” to any host in the network in terms of delay and packet loss. Results may not be as accurate as the SLA statistics supplied by other Ether.Genius / Ether.Sync / Ether.Giga tests but Ping tests are fast, virtually supported by any IP network element and at least they are a good way check end-to-end network connectivity before running a more sophisticated test.

8.1.1. Internet Control Message Protocol

Ping is an application of the Internet Control Message Protocol (ICMP) Echo request and Echo reply messages. The IP protocol alone is unable to monitor whether the packets arrive to final destination. Moreover, it does not provide any error reporting when routing and forwarding anomalies occur. This task is left to the ICMP protocol.

ICMP is a network layer Internet protocol that provides mechanisms to report errors and other information regarding IP packet processing back to the source. It is used for error reporting and analysis, transferring messages from routers and stations, and for reporting network configuration and performance problems.

ICMP generates several kinds of useful messages, including *Destination Unreachable*, *Echo Request* and *Echo Reply*, *Redirect*, *Time Exceeded*, and *Router Advertisement* and *Router Solicitation*. The ICMP functionality includes: Report network errors, Congestion indication, Troubleshooting assistance, Announce packet time-outs when TTL field is set to zero.

8.1.2. Test Configuration

The setup of a Ping test is much more simple than other tests. It does not require any special remote device like a traffic reflector to work and it works in virtually any network. The steps to follow to configure the IP ping are:

1. Make sure that the Port A of your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
3. Configure *IP endpoint* with the help of the *Mode* setting.
4. Configure *Ping/Traceroute* to *Ping*.
5. From the *Home* panel, go to *Setup*
   The test port settings panel is displayed.
6. Select *Port A* to enter in the port A specific configuration.
7. Configure the *Local profile* (See section 2.3) either by means the DHCP protocol or by hand.
8. Go to *Ping/Traceroute*
9. Enter the destination IPv4 address by using the *Destination IPv4 address from*, *Destination IPv4 address* and *Destination host name*.
10. Configure the *Timeout*, *Interval*, *ICMP packet size* and *TTL* parameters for the test.
11. Run the Ping test with the help of the *RUN* fixed button.

Table 8.1: IP Ping Settings

Setting	Description
Destination IPv4 address from	Establishes the origin of the destination IPv4 address for the current stream. There are two different settings available for configuration: • *Manual*: The destination address is set to the value configured in *Destination IPv4 address*. • *Host name*: Uses the Domain Name Service (DNS) to set the destination IP address by using descriptive alphanumeric strings. The DNS mechanism requires intervention of at least one DNS server. The DNS server IP address has to be configured in the local port profile either statically or by means DHCP.

Table 8.1: IP Ping Settings

Setting	Description
Destination IPv4 address	Destination IPv4 address carried by the packets generated in the current stream if *Destination IPv4 address from* is set to *Manual*. The address is entered in decimal, four-dotted format. Any address between 0.0.0.0 and 255.255.255.255 is admitted as a destination IPv4 address.
Destination IPv4 address (DNS)	Destination IPv4 address carried by the packets generated in the current stream if *Destination IPv4 address from* is set to *Host name*. This is a read only field that it cannot be edited directly. It displays the result of the DNS name resolution carried out with the host name configured in *Destination host name.*
Destination host name	Domain name to be used as a destination if *Destination type* is set to *Domain name*. Unlike IP addresses, domain names are easy-to-remember alphanumeric strings but they have to be translated to IP addresses before any packet can be sent to the destination. The translation process requires the intervention of at least one DNS server. The DNS server IP address has to be configured in the local port profile either statically or by means DHCP.
Timeout	Time the receiver waits for an ICMP Echo replay for each ICMP Echo request it generates. No new Echo request is issued until either the previous reply is received or the timeout period finishes without any Echo reply reception. The default value for *Timeout* is 5.0 seconds.
Interval	Separation between two consecutive ICMP transmissions. The effective separation may be longer than the value configured in this field if the time it takes to receive the corresponding ICMP Echo reply is longer than the *Interval*. The default value for *Interval* is 1.0 seconds
ICMP packet size	Packet size used in transmitted ICMP Echo requests. In order to obtain the total frame size it is necessary to add the ICMP header length (8 bytes), IPv4 datagram header length (20 bytes) and the Ethernet (DIX) frame header and trailer (18 bytes). The default value is 56 bytes that is equivalent to a frame size of 102 bytes (56 bytes + 8 bytes + 20 bytes + 18 bytes).

Table 8.1: IP Ping Settings

Setting	Description
TTL	Initial Time To Live value configured in the packets transmitted in the current stream. The TTL is decreased by one unit each time it lefts a network node. If the value reaches zero, then the packet is discarded. The TTL is then a measure of the number of nodes the packet is allowed to transverse before reaching its destination.

8.1.3. Result verification

Ping results are presented in real time within a dedicated Ether.Genius / Ether.Sync / Ether.Giga panel. To display the Ping results follow these steps:

Table 8.2: IP Ping Results

Metric	Description
Requests sent	Number of ICMP Echo request messages sent from the beginning of the test.
Replies received	Number of ICMP Echo reply messages received from the destination. If the number of replies received is smaller than the request send, then it can be concluded than the network has dropped some packets. On the other hand, if the replies number is higher than the request count, then the network is probably duplicating packets somewhere.
Replies lost	Is the count of ICMP Echo reply messages lost from the beginning of the Ping test.
Packet loss	This metric is the ratio between the ICMP Echo request messages sent and the ICMP Echo reply messages lost.
Minimum delay	Minimum delay between the Echo request message transmission and the corresponding Echo reply reception event , computed over all the available Echo request / reply pairs. The Ping minimum delay can be considered an estimate of the minimum round trip delay between source and destination but the Ping delay may include ICMP protocol processing delays in intermediate router and in the destination.

Table 8.2: IP Ping Results

Metric	Description
Maximum delay	Maximum delay between the Echo request message transmission and the corresponding Echo reply reception event computed over all the available Echo request / reply pairs,. The maximum delay figure is subject to the same non-zero processing delay uncertainties than the *Minimum delay*.
Average delay	Mean delay between Echo request transmission and Echo reply reception events computed over all the available Echo request / reply pairs. The mean ping delay is subject to the same processing delay uncertainties than the *Minimum delay*.

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select *Port A* to enter in the port specific results.
3. Select *Ping*.
4. Check the *Requests sent*, *Replies received*, *Replies lost*, *Packet loss*, *Minimum delay*, *Maximum delay* and *Average delay* statistics.

Figure 8.1: Ether.Genius / Ether.Sync / Ether.Giga Ping results panel.

8.2. Traceroute

Traceroute can be defined as an extended Ping test that traces intermediate network elements between the source and destination. These elements are identified by their IP addresses. For this reason, Traceroute is often used to check the path the packets

follow when they are transmitted to the network. Ether.Genius / Ether.Sync / Ether.Giga support two different implementations of the Traceroute application based on different probe packets:

- *ICMP Traceroute*: The test equipment uses ICMP Echo request messages with increasing TTL values to identify the nodes in the test path. Each node decreases the TTL value in one unit before forwarding the test packet. If the TTL reaches the value of 0, then the hop replies with an ICMP Time to live exceeded message to the transmission source.
- *UDP Traceroute*: It works in a similar way than the UDP Traceroute but in this case the probe packet is a regular UDP packet directed to an arbitrary port rather than an ICMP Echo request message.

It is important to notice than the some network elements located in the Traceroute test part will never be detected by this test as in is required that these elements implement the IP protocol stack reply to be able to reply to the probe packets. Ethernet switches, broadband modes operating in bridged modes and media converters fall in this category.

8.2.1. Test Configuration

Connection setup is the same for Traceroute and Ping tests. Traceroute also shares the simple and quick setup procedure with the Ping test. To configure traceroute follow these steps:

1. Make sure that the Port A of your tester is connected to the network. The physical layer must be up and working (See section 4.1.1).
2. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
3. Configure *IP endpoint* with the help of the *Mode* setting.
4. Configure *Ping/Traceroute* to *Traceroute*
5. From the *Home* panel, go to *Setup*
   The test port settings panel is displayed.
6. Select either Port A to enter in the port specific configuration.
7. Configure the *Local profile* (See section 2.3) either by means the DHCP protocol or by hand.
8. Go to *Ping/Traceroute*
9. Enter the destination IPv4 address by using the *Destination IPv4 address from*, *Destination IPv4 address* and *Destination host name*.
10. Configure the *Timeout*, *Interval*, *Max. number of hops*, *Number of packets/hop*, *Traceroute protocol* and *UDP port* (UDP Traceroute only) parameters for the test.

11. Run the Traceroute test with the help of the *RUN* fixed button.

Table 8.3: Traceroute Settings

Setting	Description
Destination IPv4 address from	This field has the same meaning than the *Destination IPv4 address* from setting used by the Ping test (See section 8.1.2).
Destination IPv4 address	This field has the same meaning than the *Destination IPv4 address* setting used by the Ping test (See section 8.1.2).
Destination IPv4 address (DNS)	This field has the same meaning than the *Destination IPv4 address (DNS)* setting used by the Ping test (See section 8.1.2).
Destination host name	This field has the same meaning than the *Destination host name* setting used by the Ping test (See section 8.1.2).
Timeout	Time the receiver waits for an ICMP Time-to-live exceeded reply for each UDP request (UDP Traceroute) or ICMP Echo request (ICMP Traceroute). No new UDP / ICMP Echo request is issued until either the previous reply is received or the timeout period finishes without any Echo reply reception. The default value for *Timeout* is 5.0 seconds.
Interval	Separation between two consecutive ICMP / UDP transmissions. The effective separation may be longer than the value configured in this field if the time it takes to receive the corresponding ICMP Port unreachable / Time-to-live exceeded is longer than the *Interval*. The default value for *Interval* is 1.0 seconds
Max. Number of hops	Maximum length of the path to be analysed with Traceroute expressed in the number of hops (routers, hosts) it contains. The test fails to reach the end point if the path to be tested contains more hops than the number configured in this field. The default number of hops is 30.
Number of packets/ hop	Number of test packets directed to each hop in the test path. The default is 1 hop but this value could be increased to reduce the variability of of delay statistics or to get more accurate packet loss statistics.

Table 8.3: Traceroute Settings

Setting	Description
Traceroute protocol	It is one of ICMP or UDP. These protocols define totally different Traceroute probe packets • ICMP: Uses ICMP Echo request messages with increasing TTL values as probe packets. • UDP: Uses UDP requests with increasing TTL values directed to a user configurable port as probe packets.
UDP port	Configures the destination UDP port used in UDP Traceroute tests. This setting is not required for ICMP Traceroute tests.

8.2.2. Result verification

Traceroute results are presented in real time within a dedicated Ether.Genius / Ether.Sync / Ether.Giga panel. To display the Traceroute results follow these steps:

Figure 8.2: Ether.Genius / Ether.Sync / Ether.Giga Traceroute results panel.

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select *Port A* to enter in the port specific results.
3. Select *Traceroute*.

4. Check the *Node* list and the *Pkts*, *Min.(ms)*, *Max.(ms)* and *Avg.(ms)* statistics for each node in the list.

Table 8.4: IP Ping Results

Result	Description
Node	IP address corresponding to a network node detected by the Traceroute test. The are displayed in a list following the same order than they have been detected by the probe packets.
Pkts	Number of ICMP Time to live exceeded messages received from a particular node in the path.
Min.(ms)	Minimum delay between the probe packet transmission and the corresponding ICMP reply reception event computed over all request / reply pairs for a particular node The Traceroute minimum delay can be considered an estimate of the minimum round trip delay between the source and the node but the Traceroute delay may include protocol processing delays in intermediate network elements.
Max.(ms)	Maximum delay between the probe packet transmission and the corresponding ICMP reply reception event computed over all request / reply pairs for a particular node. The maximum delay figure is subject to the same non-zero processing delay uncertainties than the *Min (ms)* .delay
Avg.(ms)	Maximum delay between the probe packet transmission and the corresponding ICMP reply reception event computed over all request / reply pairs for a particular node. The mean ping delay is subject to the same processing delay uncertainties than the *Min.(ms)* delay.

Chapter 9
IEEE 1588 Analysis

Ether.Genius and Ether.Sync may optionally supply IEEE 1588 version 2, also known as *Packet Time Protocol (PTP)*, emulation and analysis features. IEEE 1588 testing is not available for Ether.Giga.

Some of the potential applications for the IEEE 1588 emulation and testing capabilities included by Ether.Genius / Ether.Sync are:

- PTP master emulation synchronized with an external clock source based on 2048 kb/s, 2048 kHz, 1 pps, Synchronous Ethernet and other interfaces.
- PTP slave emulation with extended performance measurements from the associated master and network, including master identity and status, PDV analysis and frequency / phase offsets.
- Passive monitoring testing of the communication between PTP master and slave. Connection of the passive monitor could be done in end-point or pass-through modes. Supported monitoring tests are related with master identity and status, PDV analysis and frequency offset.

9.1. Ethernet Synchronization with IEEE 1588

The Precision Time Protocol (PTP), included in IEEE standard 1588 was originally designed to provide timing for critical industrial automation. With the 2008 version of this standard (IEEE 1588v2), PTP overcomes effects of latency and jitter through chains of Ethernet switches, providing accuracy in the nanosecond range.

9.1.1. Precedents: IP Synchronization with NTP

The Network Time Protocol (NTP), is one of the oldest protocols still in use and it is available in two flavours: the full version and Simple NTP (SNTP), a subset of NTP.

The latest version of NTP, version 4 (NTPv4) can usually maintain time to within 1-20 ms using traditional software-interrupt based solutions over the public Internet and can achieve accuracies of microseconds or better in LANs under ideal conditions. NTP has been the most common and arguably the most popular synchronization solution,

because it performs well over LANs and WANs and at the same time it is inexpensive, requiring very little hardware.

NTP should be able to deliver accuracy of 1-2 ms on a LAN and 1-20 ms on a WAN. However, protocol performance is far from guaranteed largely because of variable delays added by switches and routers.

9.1.2. PTP Protocol Details

PTP only requires a central Grandmaster clock and low-cost PTP slave clock sites. Master and slave network devices are kept synchronized by the transmission of timestamps transmitted within the PTP messages.

Depending on how many ports has a network clock, it is referred by the IEEE 1588 standard as an *Ordinary Clock* (single port device) or a *Boundary Clock* (multi port device). The version 2 standard also defines the concept of Transparent Clocks that improve timing accuracy when the protocol is run in network paths which contain intermediate switches.

Table 9.1: IEEE 1588v2 Device Description

Device	Description
Ordinary Clock	A single port device that can be a master or slave clock.
Boundary Clock	A multi port device that can be a master or slave clock.
End-to-end Transparent Clock	A multi port device that is not a master or slave clock but a bridge between the two. Forwards and corrects all PTP messages.
Peer-to-peer Transparent Clock	A multi port device that is not a master or slave clock but abridge between the two. Forwards and corrects Sync and Follow-up messages only.
Management Node	A device that configures and monitors clocks.

The normal execution of the PTP has two phases:

1. *Master-Slave hierarchy establishment*. Ordinary and boundary clocks decide which port has the master or slave role in each link with the help of the Best Master Clock (BMC) algorithm. The required data required for operation of the BMC is supplied by special *Announce* messages generated periodically by Ordinary and Boundary Clocks.
2. *Clock synchronization*. Slave clocks may have a positive or negative offsets when compared with their masters and latency from masters to slaves is also unknown. PTP devices start a procedure to compute latencies and offsets. These parameters will be used to adjust timing in slave devices.

9.1.3. The Synchronization Mechanism

Once the master and slave hierarchies have been established, by observing the clock property information contained in *Announce* messages sent by PTP devices, the synchronization process starts.

Before synchronization between the master and the slave clock has been achieved, it may exist an offset between both clocks. This offset is computed with the help of the *Sync* message. *Sync* messages are sent periodically (usually once every one second) by the master to upgrade offset information in the slave. *Sync* messages may carry an accurate timestamp indicating the departure time of the own message but this requires expensive timestamping hardware which may not be available. To avoid expensive hardware *Follow_Up* messages can be used. *Follow_Up* messages carry timestamps for a previous *Sync* message allowing a more relaxed timestamping procedure and cheaper hardware.

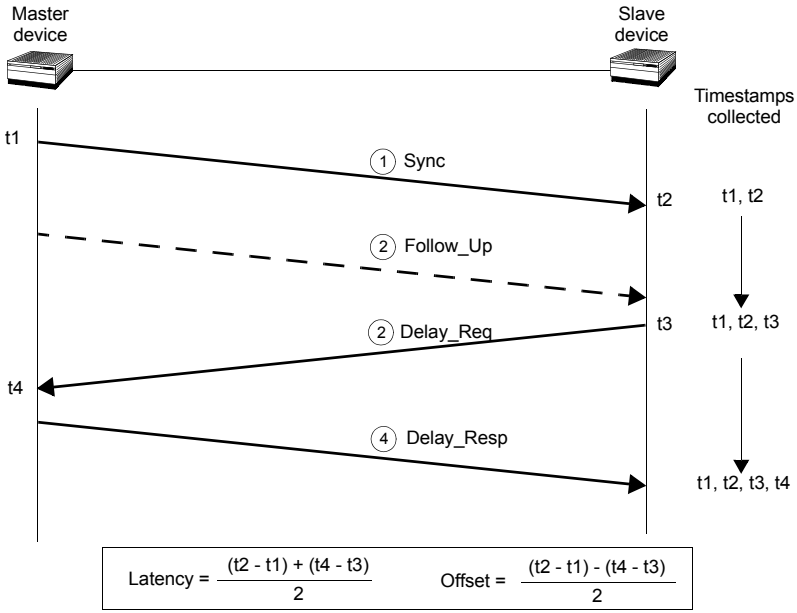

Figure 9.1: Sync and Delay request-response mechanisms used by the PTP. The basic parameters of Latency and Offset are computed from the t1, t2, t3 and t4 timestamps.

The *Sync* mechanism, however, does not take into account propagation time of *Sync* messages through the network. For this reason, the slave may request a latency measurement with a *Delay Req* message. Masters reply to a *Delay Req* with *Delay_Resp* message. The timestamps the slave get from the Delay Request-Response mechanism are used to correct *Sync* with a more accurate time estimation.

The most difficult challenge of PTP is operation though chains of Ethernet switches. Most switches store packets in local memory while the MAC address table is searched and the cyclic redundancy field of the packet is checked before it is sent out on the appropriate port/s. This process introduces variations in the time latency of packet forwarding and damages accuracy of the PTP protocol. Version 1 of the PTP protocol deal with this problem by implementing Boundary Clocks within the switches. Version 2 uses the more advanced concept of Transparent Clock to deal with the same problem.

Transparent clocks do not participate in the master-slave hierarchy but they process PTP messages by adding special correction fields within the message based on their own estimations of packet residence times in the device. There are two different kinds of Transparent Clocks depending on whether they use the Peer-delay mechanism or the Delay request-response mechanism to compute the propagation delay between master and slave. The Peer-delay mechanism is more sophisticated and it has more pre-requisites than the Delay-request response mechanism but the Peer-delay mechanism is more scalable. End-to-end Transparent clocks add residence time compensations to all the PTP event messages but Peer-to-peer Transparent Clocks are only required to compensate *Sync* and *Follow-up* messages because they do not use the *Delay_Req* and *Delay_Resp* (Delay request-response mechanism) and they do not need to forward the *Pdelay_Req* and *Pdelay_resp* messages (Peer delay mechanism).

9.1.4. Protocol Encapsulation

PTP messages can be carried over a large family of protocols including IPv4, IPv6, IEEE 802.3 Ethernet, DeviceNET, ControlNET and IEC 61158 Type 10. The most important encapsulations are the IP and Ethernet variations (see Figure 9.2:).

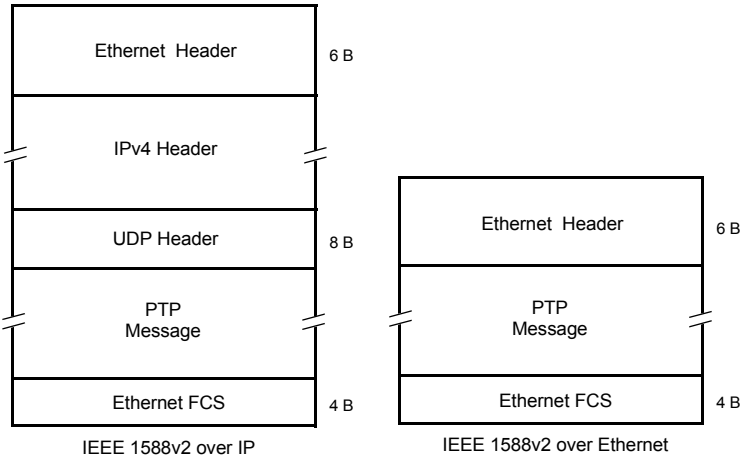

Figure 9.2: IP and Ethernet encapsulations for PTP messages.

IEEE 1588 messages encapsulated in Ethernet frames use the special 0x88f7 Ethertype, specifically reserved for this purpose. Messages associated with the peer-delay mechanism carry the 01:1B:19:00:00:00 multicast destination address while the remaining messages carry the 01:80:C2:00:00:0E multicast address. On the other hand, messages encapsulated in IPv4 datagrams, use the UDP protocol for transport and the destination port is 319 (event messages) or 320 (all other messages). Multicast addresses to be used when the encapsulation is set to IPv4 are the 224.0.0.107 (peer delay mechanism messages) and the 224.0.1.129 (all other messages).

9.2.IEEE 1588 Master and Slave Emulation

The Ether.Genius and Ether.Sync testers can be configured to behave as different kinds of PTP entities, including PTP master and slave clocks. Theoretically, there is no difference between any standard PTP clock and Ether.Genius / Ether.Sync but the testers supply measurement results that are useful to qualify the stability of a PTP timing source, the network performance when delivering different IEEE 1588 messages and the ability of PTP slaves to recover an accurate timing from the messages they receive from their masters:

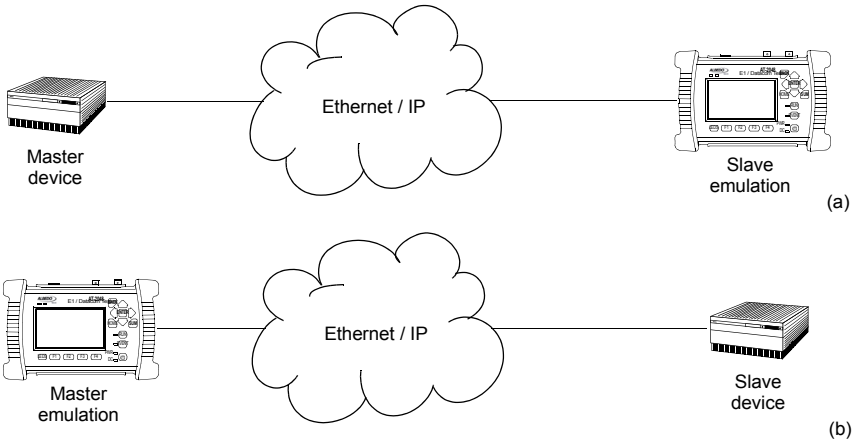

Figure 9.3: IEEE 1588 master and slave emulation modes: (a) Connection diagram corresponding to the PTP slave emulation mode, (b) connection corresponding to the master emulation mode.

The IEEE 1588 test results provided by Ether.Genius and Ether.Sync can be classified in four different families:

- *Protocol state*: Not specifically test results. The protocol state contains details received from the master or the grandmaster clock (or the own test equipment when configured in master emulation mode).

- *Message statistics*: Includes message counts classified by their type (*Sync*, *Delay Request*, *Follow up*,...).
- *Delay statistics*: The *Packet Delay Variation* (PDV) and the *Packet Total Delay* (PTD) are some of the test results included in this family. Among different things, the delay statistics are useful to measure fast variations in the network transmission conditions affecting the timing accuracy of PTP slave devices.
- *Slave clock status*. When the test unit is configured in slave emulation mode, it tracks the frequency and phase of a PTP master. Due to phase / frequency instability in the master clock, or impairments caused by transmission through the packet switched network, the slave clock has to continuously adjust its frequency and phase. This test result family quantifies the difference between the local (slave) clock and the estimated frequency / phase received from the master.

The exact test results available in any case depend on the current operation mode and on the test port. Port B contains a reduced set of PTP test results (*message statistics*) and it is unable to generate any message on its own. For this reason, for a full PTP analysis in master or slave emulation mode, it is recommended to connect the equipment to the network through the port B. Details about which results are available in each mode are provided in the following sections.

9.2.1. Configuration of the IEEE 1588 Master and Slave

When configuring Ether.Genius or Ether.Sync in master or slave emulation modes, it has to be taken into account that only the test Port A contains a full implementation of the protocol. Port B can be used to collect message statistics but transmission and decoding of PTP messages required to supply synchronization to remote entities (master emulation mode) or to synchronize the own clock with a remote entity (slave mode) are available for Port B only.

The IEEE 1588 slave and master clock emulation is compatible with the *Ethernet endpoint* and *IP endpoint* operation modes (See section 2.1). Configuration is slightly different in each case. If you are operating in *IP endpoint* mode you need to configure your local profile before you can run a PTP test. To do that, the required steps are:

1. From the *Home* panel, go to *Setup*
   The test port settings panel is displayed.
2. Select *Port A* to enter in the port specific configuration.
3. Configure the *Local profile* (See section 2.3) either by means the DHCP protocol or by hand.

The previous steps are not required if the equipment is operating in Ethernet endpoint mode. To run a PTP test in (*Ethernet endpoint* mode or *IP endpoint* modes) follow these steps:

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Go to *PTP (IEEE 1588)*.

3. Configure the equipment to become an active IEEE 1588 entity by configuring *PTP test* to *PTP clock emulation*.
   A label with the text PTP is displayed in the top notification area.
4. Configure *Clock emulation*, to *Master*, *Slave* or *Auto*, depending on your preferences.
5. Set the transport protocol to either *Ethernet* or *UDP*.
   Note: If you are operating in Ethernet endpoint mode, the *UDP* transport protocol is not allowed.
6. Configure the *Domain* to the right value for your network. The default value in most networks is 0.
7. Configure the *Clock class*, *Custom clock class*, *Priority 1* and *Priority 2* fields to the right values. These values determine the PTP role (master or slave) to be played by *Ether.Genius* or *Ether.Sync* when *Clock emulation* parameter is set to *Auto*.

Table 9.2: IEEE 1588 Settings

Setting	Description
PTP test	Configures the IEEE 1588 test to be executed. These are the options currently available for this setting: • *None*: Disables all PTP generation and analysis. • *PTP clock emulation*: Choose this mode if you want the test unit to behave as a PTP slave or master clock. The unit becomes a new PTP entity in your network it will generate and receive IEEE 1588 messages in the same way than any other master / slave clock installed in your network once it is properly configured. • *Passive monitor*: This is the right mode to verify PTP communications between a PTP master and slave without disturbing the transmission channel and the PTP entities.
Clock emulation	This setting enables the user to configure the role of the test unit within the PTP network. It is one of the following: • *Master*: The equipment is forced to assume the role of a PTP master clock within the network. • *Slave*: The equipment is forced to assume the role of a PTP slave clock within the network. • *Auto*: The equipment behaves as a PTP ordinary clock. It becomes a master or slave depending on the result of the BMC algorithm.

Table 9.2: IEEE 1588 Settings

Setting	Description
Domain	Configures the PTP domain for the test unit. The PTP domain is identified by a number between 0 and 255. The equipment is allowed to exchange PTP information only with clocks within the same domain. The unit ignores all messages received from other domains (these messages are classified as *Domain mismatches*). All PTP equipments from other domains will probably ignore all messages from the tester.
Clock class	Sets the clock class as defined in IEEE 1588-2008 or the clock-source quality-level defined in ITU-T G.781. The Clock class defines traceability of the time or frequency distributed by the master clock and it is one of the parameters used to establish the master / slave hierarchy with the help of the BMC. This field is available for configuration only if *Clock emulation* has been set to *Auto*.
Custom clock class	The *Custom clock class* sets a custom clock class in numeric format when *Clock class* has been configured to *Custom*. Any *Custom clock class* between 0 and 255 is allowed here.
Priority 1	This is a numeric parameter used by the BMC to establish the master / slave hierarchy with the help of the BMC algorithm. The range of accepted values of *Priority 1* is between 0 and 255. It is more likely that the test unit becomes a PTP master during the BMC if *Priority 1* is configured to an small value. This field is available for configuration only if *Clock emulation* has been set to *Auto*.
Priority 2	This is a numeric parameter used by the BMC to establish the master / slave hierarchy with the help of the BMC algorithm. The range of accepted values of *Priority 2* is between 0 and 255. It is more likely that the test unit becomes a PTP master during the BMC if *Priority 2* is configured to an small value. This field is available for configuration only if *Clock emulation* has been set to *Auto*.

Table 9.2: IEEE 1588 Settings

Setting	Description
Transport protocol	Configures the encapsulation protocol for the PTP messages. It is one of the following: • *Ethernet*: PTP messages are transmitted and received through an IEEE 802.3 / Ethernet encapsulation as specified in IEEE 1588-2008 Annex F. • *UDP*: PTP messages are encapsulated in UDP over IPv4 frames as specified by IEEE 1588-2088 Annex D.
Master identity	Menu that contains the menu items necessary to set the master identity either as a MAC address, IPv4 address or host name. You can also leave the equipment to choose the master identity for you if you set the *Source of identity* subfield to *Auto*. Configuring the master identity makes sense only you have configured the equipment as a *Passive monitor*.
Message timing	This menu contains menu items necessary to configure the timing associated to the transmission of several message types when the equipment is configured in clock emulation mode. With the help of the *Message timing* menu you can configure the following parameters: *Sync TX interval*, *Announce TX interval*, *Announce RX timeout (#msgs)*, *Delay Request TX interval*.

8. Optionally, configure the timing of the different messages associated to PTP from the Message timing menu.

When Ether.Genius or Ether.Sync is configured in endpoint mode and PTP entity emulation is enabled, the equipment can still be configured as a multistream traffic generator (See section 4.1, See section 4.2). This feature is useful to verify the performance of the PTP protocol with different traffic loads.

9.2.2. Protocol State

The basic PTP results are available in the *Protocol State* panel. All the information contained in this panel is collected from *Announce* messages received from remote PTP entities or is an indication of the state of the internal PTP synchronization machine.

The *Protocol State* results are permanent, it is not required to start a test (*RUN* button) to display the information from this panel. The information is updated in real-time as new changes in the protocol state are registered by the test unit. To display the protocol state panel follow these steps:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.

2. Select either *Port A* to enter in the port A specific results.

Table 9.3: Protocol State Results

Result	Description
Port state	Displays one of the port states defined for the BMC algorithm in standard IEEE 1588-2008. These states determine which is the current role of the network equipment within the synchronization network (master or slave) or inform about special conditions or problems within the BMC algorithm. The Port State is set to one of the following values: • *Initializing*: While a port is in the *Initializing* state, the port initializes its data sets, hardware, and communication facilities. The port does not place any PTP messages on its communication path. • *Faulty*: The fault state of the protocol. A port in this state does not place any PTP messages. • *Disabled*: The port does not place any messages on its communication path. All PTP received messages are ignored except for statistic compilation. • *Listening*: The port is waiting for an announce receipt timeout to expire or to receive an *Announce* message from a master. The purpose of this state is to allow orderly addition of clocks to a domain. • *Pre-master*: The port behaves in all respects as though it were in the *Master* state except that it does not place any messages on its communication path except for *Peer Delay Request*, *Peer Delay Response* or *Peer Delay Follow up*. • *Master*: The port is behaving as a master clock. • *Passive*: The port does not place any messages on its communication path except for *Peer Delay Request*, *Peer Delay Response* or *Peer Delay Follow up*. • *Uncalibrated*: One or more master ports have been detected in the domain. The appropriate master port has been selected, and the local port is preparing to synchronize to the selected master port. This is a transient state to allow initialization of synchronization servos, updating of data sets when a new master port has been selected. • *Slave*: The port is synchronizing to the selected master port.

Table 9.3: Protocol State Results

Result	Description
Master identity	EUI-64 code associated to the master clock. The EUI-64, computed as explained in standard IEEE 1588-2008, constitutes a globally unique identifier.
Grandmaster identity	EUI-64 code associated to the grandmaster clock. The EUI-64, computed as explained in standard IEEE 1588-2008, constitutes a globally unique identifier. The Grandmaster identity is different to the master identity only when the synchronization is transmitted through a chain involving several PTP master / slave relations.
Grandmaster priority 1	Priority 1 parameter configured in the grandmaster. The *priority 1* is a numeric parameter used by the BMC algorithm to establish the master / slave hierarchy. The smaller this value is, the higher priority is assigned to the clock to become master.
Grandmaster priority 2	Priority 2 parameter configured in the grandmaster. The *priority 2* is a numeric parameter used by the BMC algorithm to establish the master / slave hierarchy. The smaller this value is, the higher priority is assigned to the clock to become master.
Grandmaster clock class	Clock class associated to the grandmaster clock. The Clock class defines traceability of the time or frequency distributed by the master clock and it is one of the parameters used to establish the master / slave hierarchy with the help of the BMC.
Grandmaster clock accuracy	The grandmaster clock accuracy indicates the expected accuracy of a clock when it is the grandmaster, or in the event it becomes the grandmaster. This is a field that characterized a clock for the purpose of determining the master / slave hierarchy.
Grandmaster clock variance	Grandmaster clock statistic variance estimated as specified in IEEE 1588-2008. This field is used by the grandmaster clock to report the variability of its own internal oscillator.

Table 9.3: Protocol State Results

Result	Description
Grandmaster time source	Indicates the time source used by the grandmaster clock. The value is not used in the selection of the grandmaster clock. It is one of the following • *Atomic clock*: Any device that is based on atomic resonance for frequency and that has been calibrated against international standards for frequency and time. • *GPS*: Any device synchronize to a satellite system that distribute time and frequency tied to international standards. • *Terrestrial radio*: Any device synchronized via any of the radio distribution systems that distribute time and frequency tied to international standards • *PTP*: Any device synchronized to a PTP-based source of time external to the domain • *NTP*: Any device synchronized via the Network Time Protocol (NTP) or the Simple Network Time Protocol (SNTP) to servers that distribute time and frequency tied to international standards. • *Hand set*: Used for any device whose time has been set by means of a human interface based on observation of an international standards source of time to within the claimed clock accuracy. • *Other*: Other source of time and / or frequency not covered by other values. • *Internal oscillator*: Any device whose frequency is not based on atomic resonance nor calibrated against international standards for frequency, and whose time is based on a free running oscillator with epoch determined in an arbitrary or unknown manner.

3. Enter in *PTP* to display results about the PTP protocol.
4. Go to *Protocol state* and check *Port state*, *Master identity*, *Grandmaster identity*, *Grandmaster priority 1*, *Grandmaster priority 2*, *Grandmaster clock class*, *Grandmaster clock accuracy*, *Grandmaster clock variance*, *Grandmaster time source*.

9.2.3. Message Statistics

The *Message statistics* panel include counts of each PTP message type defined in IEEE 1588-2008. This result panel includes statistics both about received and transmitted (internally generated) packets.

There is one Message statistics panel for each test port. In case, that protocol emulation or delay / jitter tests are not required, Port B can be used for message statistics compilation. The procedure to display the message statist is is as follows:

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select either *Port A or Port B* to enter in the port specific results.
3. Enter in *PTP* to display results about the PTP protocol.
4. Go to *Message statistics*.
5. Press *RUN* to start a new test and check the transmitted (TX) and received (RX) values of *Sync*, *Delay Request*, *Delay Response*, *Peer Delay Request*, *Peer Delay Response*, *Follow-up*, *Peer Delay*, *Follow-up*, *Announce*, *Signaling*, *Management* and *Domain mismatch*,

Table 9.4: IEEE 1588 Message Statistics

Result	Description
Sync	Accounts for the number of received (RX) and transmitted (TX) *Sync* messages. *Sync* messages are generated by PTP master clocks and carry information about their clock frequency and phase. To achieve accurate synchronization the PTP slave needs to complement the information obtained from the Sync messages either with the Delay request-response or the Peer delay mechanisms.
Delay Request	Number of received (RX) and transmitted (TX) *Delay Request* messages. The *Delay Request* message is part of the *Delay request-response* mechanism used to measure the two-way latency between master and slave. The *Delay Request* message is always generated by the PTP slave.
Delay Response	Number of received (RX) and transmitted (TX) *Delay Response* messages. The *Delay Response* message is part of the *Delay request-response* mechanism used to measure the two-way latency between master and slave. The *Delay Request* message is always generated by the PTP master as a reply to a *Delay request* received from a slave.

Table 9.4: IEEE 1588 Message Statistics

Result	Description
Peer Delay Request	Computed number of received (RX) and transmitted (TX) *Peer Delay Request* messages. The *Peer Delay Request* message is part of the *Peer Delay* mechanism used to measure the two-way latency between master and slave.
Peer Delay Response	Number of received (RX) and transmitted (TX) *Peer Delay Response* messages. The *Delay Response* message is part of the *Delay request-response* mechanism used to measure the two-way latency between master and slave. The *Delay Request* message is always generated by the PTP master as a reply to a *Delay request* received from a slave.
Follow-up	Number of received (RX) and transmitted (TX) *Follow-up* messages. *Follow up* packets are generated by some PTP masters when for some reason *Sync* messages cannot carry accurate timestamps. In this case the a *Follow up* message carrying the accurate timestamp is transmitted immediately after every *Sync* packet.
Peer Delay Follow-up	Number of received (RX) and transmitted (TX) *Peer Delay Follow-up* messages. *Peer Delay Follow-up* packets are generated by some PTP devices using the *Peer delay* mechanism masters when for some reason *Sync* messages cannot carry accurate timestamps. In this case the a *Follow up* message carrying the accurate timestamp is transmitted immediately after every *Sync* packet.
Announce	Number of received (RX) and transmitted (TX) *Announce* messages. Announce messages are generated by PTP ordinary and boundary clocks in order to establish or modify a master / slave hierarchy by means the *Best Master Clock* (BMC) algorithm.
Signaling	Number of received (RX) and transmitted (TX) *Signalling* messages. *Signaling* messages are generated by PTP entities to negotiate some optional features like unicast transmission.

Table 9.4: IEEE 1588 Message Statistics

Result	Description
Management	Number of received (RX) and transmitted (TX) *Management* messages. Management messages are used to connect the PTP entities to the management system and enable operation and maintenance of the synchronization network.
Domain mismatch	Total count of received IEEE 1588 messages no corresponding to the currently configured PTP domain.

9.2.4. Delay Statistics

The *Delay statistics* panel is probably one of the most important PTP results panel as it valuable quantitative information to assess the performance of the transmission network. Like it happens with most of the PTP results, the *Delay statistics* are available only for Port A.

Delay statistics does not depend on the operation mode (*IP endpoint* or *Ethernet endpoint*) but they do on the clock role. For example, in master emulation mode, *Sync* messages are ignored and therefore all metrics related with this message are not available. The PTP delay statistics computed by Ether.Genius and Ether.Giga are compensated with the help of the correction field included in PTP event messages. The objective of the delay compensation is to evaluate the the transmission performance achieved with switches behaving as transparent clocks. The procedure to display the delay statistics associated to PTP messages is as follows:

Table 9.5: IEEE 1588 Delay Statistics

Metric	Description
Sync PTD (current)	Last calculated value of the Packet Total Delay (PTD) experienced by IEEE 1588 *Sync* messages when they travel from the master to the test unit (configured in slave emulation mode). The result is expressed in microseconds. The Sync *PTD (current)* metric is compensated in delay which means that corrections carried out by intermediate IEEE 1588 transparent clocks are taken into account to compute the end result. This result is available only if the equipment is operating in slave emulation mode.

Table 9.5: IEEE 1588 Delay Statistics

Metric	Description
Sync PTD (minimum)	Minimum value of *Sync PTD (current)* registered from the beginning of the test. This result is available only if the equipment is operating in slave emulation mode.
Sync PTD (maximum)	Maximum value of *Sync PTD (current)* registered from the beginning of the test. This result is available only if the equipment is operating in slave emulation mode.
Sync PTD (average)	Mean value of the PTD computed over all *Sync* messages received from the beginning of the test and expressed in microseconds. This result is available only if the equipment is operating in slave emulation mode.
Sync PTD (std. dev.)	Standard deviation of the PTD computed over all Sync messages received from the beginning of the test and expressed in microseconds. Note that even if it is an statistic corresponding to the PTD, the standard deviation is a quantity related with how the PTD varies over the collected samples. In fact, the *Sync PTD (std. dev.)* does not depend on the absolute delay from the master and it is computed both in slave emulation and passive monitoring modes.
Sync PTD (range)	Difference between the *Sync PTD (maximum)* and *Sync PTD (minimum)*. This is an Sync PTD statistic but like it happens with the *Sync PTD (std. dev.)* is more related with delay variation than with absolute delay. It could be that the *Sync PTD (range)* is known even if *Sync PTD (maximum)* and *Sync PTD (minimum)* are not. For this reason, the *Sync PTD (range)* statistic is computed both in slave emulation and passive monitoring modes.
Sync PDV (current)	Current value of t the Packet Delay Variation (PDV) computed as per RFC 3393 and RFC 1889. Delay variation is computed over consecutively transmitted packets. The instantaneous value is smoothed with the function defined in RFC 1889 before being displayed. The *Sync PDV (current)* is computed both in slave emulation and passive monitoring modes.

Table 9.5: IEEE 1588 Delay Statistics

Metric	Description
Sync PDV (maximum)	Maximum value of *Sync PDV (current)* registered from the beginning of the test. The *Sync PDV (maximum)* is computed both in slave emulation and passive monitoring modes.
Sync PDV (average)	Mean value of all the Sync PDV values computed from the beginning of the test. Each individual delay variation is evaluated as the absolute value of the FTD associated to a given frame minus the FTD associated to the frame transmitted next. All possible consecutive frame transmission events are taken into account for the calculation of this performance metric. The only exception to this rule is if one or both frames are lost. The *Sync PDV (average)* is computed both in slave emulation and passive monitoring modes.
Delay req. PTD (current)	Last calculated value of the PTD experienced by IEEE 1588 *Delay req.* messages when they travel from the test unit (configured in slave emulation mode) to the master or when they are transmitted from the slave to the test unit (configured in master emulation mode). The result is expressed in microseconds.
Delay req. PTD (minimum)	Minimum value of *Delay req. PTD (current)* registered from the beginning of the test. The *Delay req. PDV (minimum)* is computed both in slave emulation and master emulation operation modes.
Delay req. PTD (average)	Mean value of the PTD computed over all *Delay req.* messages received (master emulation mode) or transmitted (slave emulation mode) from the beginning of the test and expressed in microseconds.
Delay req. PTD (std. dev.)	Standard deviation of the PTD computed over all Delay request messages received from the beginning of the test and expressed in microseconds. This metric is computed when the equipment is operating in IEEE 1588 master clock emulation or IEEE 1588 slave clock emulation.
Delay req. PTD (range)	Difference between *Delay req. PTD (maximum)* and *Delay req. PTD (minimum)*. This metric is computed only in slave emulation and master emulation modes.

Table 9.5: IEEE 1588 Delay Statistics

Metric	Description
Two-way PTD (current)	This result is computed as the sum of the *Sync PTD (current)* and *Delay req. PTD (current).* It is an estimation of the time invested by a two-way IEEE 1588 packet transmission between the master and the slave (or the slave and the master). This metric is computed in slave emulation mode only.
Two-way PTD (minimum)	This result is computed as the sum of the *Sync PTD (minimum)* and *Delay req. PTD (minimum).* It is an estimation of the minimum time invested by a two-way IEEE 1588 packet transmission between the master and the slave (or the slave and the master). This metric is computed in slave emulation mode only.
Two-way PTD (average)	Metric calculated as the sum of the *Sync PTD (average)* and *Delay req. PTD (average)*. It is an estimation of the average time invested by a two-way IEEE 1588 packet transmission based on the packets collected from the beginning of the test. The *Two-way PTD (average)* is computed only in slave emulation mode.
Sync IAD (current)	Current Inter Arrival Delay (IAD) computed over the *Sync* packets received from a remote IEEE 1588 master clock. The Sync IAD is defined as the delay between to consecutively received IEEE 1588 Sync packets. The *Sync IAD (current)* metric is computed in slave emulation and passive monitor modes.
Sync IAD (average)	Mean value of the Sync IAD computed over all *Sync* messages received from the beginning of the test and expressed in microseconds. The *Sync IAD (average)* metric is computed in slave emulation and passive monitor modes.

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select *Port A* to enter in the port specific results.
3. Enter in *PTP* to display results about the PTP protocol.
4. Go to *Delay statistics*.

5. Press *RUN* to start a new test and check the values of *Sync PTD (current)*, *Sync PTD (minimum)*, *Sync PTD (maximum)*, *Sync PTD (average)*, *Sync PTD (std. dev.)*, *Sync PTD (range)*, *Sync PDV (current)*, *Sync PDV (maximum)*, *Sync PDV (average)*, *Delay req. PTD (current)*, *Delay req. PTD (minimum)*, *Delay req. PTD (maximum)*, *Delay req. PTD (average)*, *Delay req. PTD (average)*, *Delay req. PTD (std. dev.)*, *Delay req. PTD (range)*, *Two-way PTD (current)*, *Two-way PTD (minimum)*, *Two-way PTD (average)*, *Sync IAD (current)*, *Sync IAD (average)*.

9.2.5. Slave Clock Status

When Ether.Genius or Ether.Sync are operating in slave emulation mode they are required to track the frequency and phase of a signal generated by a PTP master and encoded in different types of messages. Not ideal conditions in the master, the transmission network and the slave cause transient or permanent errors withing the phase and frequency recovered by the slave equipment. These errors can be estimated on the basis of fluctuations no the received phase. The phase and frequency error are displayed in the *Slave clock status* panel. To display the phase and frequency errors, follow these steps:

Table 9.6: Slave Clock Status Metrics

Metric	Description
Frequency offset	Displays the frequency offset between currently selected IEEE 1588 master and the test unit. If the equipment is working in slave emulation or passive monitoring modes, it slowly tracks the frequency recovered from the master. In this case, under ideal transmission conditions, the *Frequency offset* should converge to an small value.
Phase offset	Phase offset between the currently selected IEEE 1588 master and the Ether.Genius / Ether.Sync test unit. If the equipment is working in slave emulation mode, it slowly tracks the phase recovered from the master (and compensates for transmission latency using either the *Delay request-response* or the *Peer delay* mechanisms). The phase offset should be close to zero under ideal or nearly ideal transmission conditions.

1. From the *Home* panel, go to *Results*,
   The test port results panel is displayed.
2. Select *Port A* to enter in the port specific results.
3. Enter in *PTP* to display results about the PTP protocol.
4. Go to *Slave clock status*.
5. Press *RUN* to start a new test and check the values of *Frequency offset* and *Phase offset*.

9.3. Passive Monitoring

Sometimes, the ability to qualify the operation of an existing PTP slave already connected to the network is more important than a full emulation of a PTP slave. A typical example happens when an slave device exhibits poor performance for unknown reasons. A quick PDV measurement with Ether.Genius / Ether.Giga could be useful to determine whether the problem is caused by the network (congestion, inappropriate switches) or the device (high sensitivity to PDV, configuration problem).

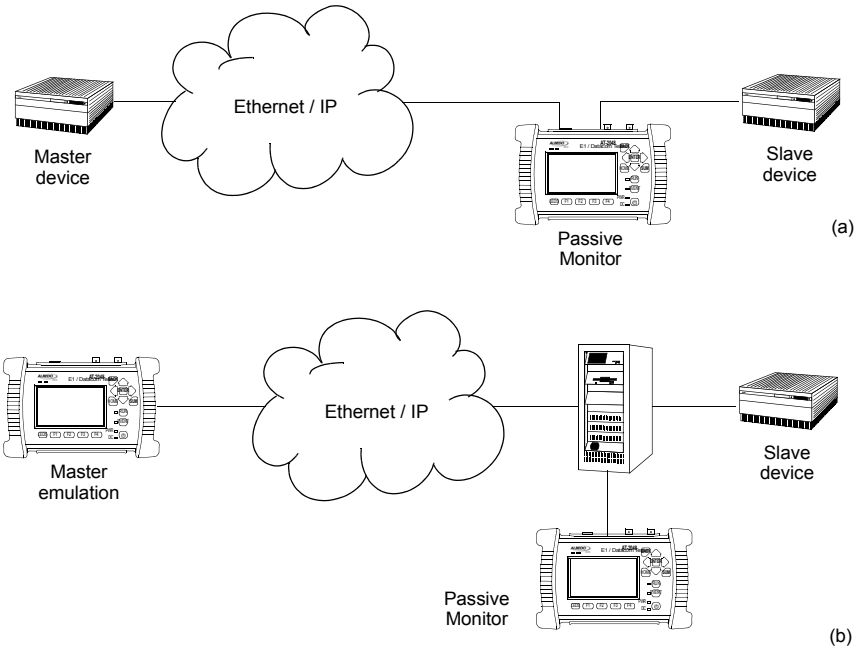

Figure 9.4: Connection of Ether.Genius / Ether.Sync to the network in PTP *Passive monitor* mode: (a) Connection in pass-through mode (IP through operation mode). (b) Connection in endpoint mode (*Ethernet endpoint* or *IP endpoint* operation modes).

The passive monitoring operation mode is defined so that it doesn't affect the network in any way. No PTP traffic is generated in this mode. Interaction is limited to analysis of the different types of messages received from other PTP master and slave clocks. Connection in Ethernet or IP endpoint is allowed when the monitoring mode is configured but the *IP through* mode is allowed as well. By connecting the equipment in pass-through mode at the output of the DUT (usually a PTP slave) it is possible to analyse the messages the DUT is receiving without any intermediate device that could be masking the real synchronization performance.

When Ether.Genius / Ether.Sync is operating in *PTP monitor* mode it cannot track the phase generated by the master because it does not estimate the path delay (Delay request-response mechanism, Peer-delay mechanism) but it can still syntonize the master frequency. The result is that all metrics related with absolute phase or delay are not available in PTP *Passive monitor* mode.

If you want to run the PTP monitor in *IP endpoint* or *IP through* modes, you will need to set your local profiles with correct IP addresses, network masks and gateways. For the specific case of the *IP through* mode, you will have to configure the local profile of both *Port A* and *Port B*. The details are as follow:

1. From the *Home* panel, go to *Setup*
   The test port settings panel is displayed.
2. Select *Port A* to enter in the port specific configuration.
3. Configure the *local profile* (See section 2.3) either by means the DHCP protocol or by hand.

If you are operation mode is *IP endpoint*, the local profile configuration finishes here. If you are working in *IP through* mode, repeat the previous steps for *Port B*. The configuration specific for the PTP protocol is detailed in the following steps:

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Go to *PTP (IEEE 1588)*.
3. Configure the equipment to become an passive IEEE 1588 monitor by configuring *PTP test* to *Passive monitor*.
   A label with the text PTP is displayed in the top notification area.
4. Set the transport protocol to either *Ethernet* or *UDP*.
   Note: If you are operating in Ethernet endpoint mode, the *UDP* transport protocol is not allowed.
5. Configure the *Domain* to the right value for your network. The default value in most networks is 0.
6. Go to *Master identity* to select a remote PTP master clock.
7. Configure Source of identity to one of the allowed values: *Auto*, *IPv4 address*, and *Host name* (UDP transport protocol) or *Auto* and *MAC address* (*Ethernet transport* protocol).
8. If you have configured *Source of identity* to *IPv4 address*, *Host name* or *MAC address*, fill the correct master identify field with the data from your master.
9. Optionally, configure the timing of the different messages associated to PTP from the Message timing menu.

Once configured, you can use the same measurement procedures described for the PTP master and slave clock emulation (See section 9.2) to get statistics and counts in your test environment. The only difference in the results is that in the *Passive monitor* mode only a subset of the statistics available for the endpoint emulation are available.

Chapter 10
Synchronous Ethernet Analysis

Ether.Genius and Ether.Sync may optionally support the ITU-T defined Synchronous Ethernet standard which enables these equipments qualify Synchronous Ethernet network equipment or to generate Ethernet signals synchronized to various timing sources, including GPS and TDM.

10.1.Introduction to Synchronous Ethernet

Synchronous Ethernet is an ITU-T standard that provides mechanisms to transfer frequency over the Ethernet physical layer, which can then be made traceable to an external source such as a network clock. As such, the Ethernet link may be used and considered part of the synchronization network.

The proposal to specify the transport of a reference clock over Ethernet links was brought by operators to ITU-T Study Group 15 in September 2004. The aim of Synchronous Ethernet is to avoid changes to the existing IEEE Ethernet, but to extend it working within its protocol definitions.

Despite being an IEEE standard, Ethernet architecture has been described in ITU-T G.8010 as a network made up of an ETH layer and a ETY layer. Put in simple terms, the ETY layer corresponds the physical layer as defined in IEEE 802.3, while the ETH layer represents the pure packet layer. Ethernet MAC frames at the ETH layer are carried as a client of the ETY layer. In OSI terminology, ETY is layer 1, ETH layer 2. Synchronous Ethernet is based on the ITU-T G.8010 description of the Ethernet architecture.

A key topic in Synchronous Ethernet is the definition of the mechanisms necessary to achieve interworking between SDH and Synchronous Ethernet equipment. These mechanisms and procedures are found fundamentally in three different recommendations: ITU-T G.8261, G.8262 and G.8264. The aspects covered there include the following:

- Extension of the synchronization network to include Ethernet as a building block (ITU-T G.8261). This enables Synchronous Ethernet network equipment to be

connected to the same synchronization network that SDH. Synchronization for SDH can be transported over Ethernet and the opposite is also true.

- The ITU-T G.8262 defines Synchronous Ethernet clocks compatible with SDH clocks. Synchronous Ethernet clocks are based on ITU-T G.813 clocks and they are defined in terms of accuracy, noise transfer, holdover performance, noise tolerance, and noise generation. These clocks are referred as Ethernet Equipment Slave clocks. While the IEEE 802.3 standard specifies Ethernet clocks to be within ±100 ppm. EECs accuracy is within ±4.6 ppm. Additionally, by timing the Ethernet clock, PRC traceability of the interface is achievable.
- ITU-T G.8264 extends the usability of the ITU-T G.707 Synchronization Status Message (SSM) by Synchronous Ethernet equipment. The SSM contain an indication of the quality level of the clock that is driving the synchronization chain. The *Ethernet Synchronization Message Channel* (ESMC) is used for propagation of the SSM through the Synchronous Ethernet network.

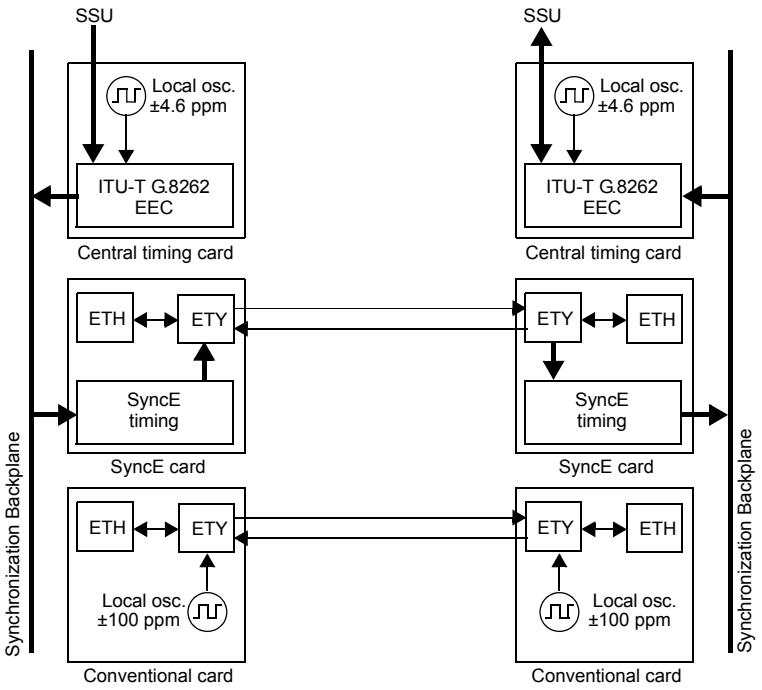

Figure 10.2 Synchronous Ethernet Architecture and comparison with conventional Ethernet

The basic difference between a conventional Ethernet and a Synchronous Ethernet network interface card is that the Synchronous Ethernet card is prepared to accept

external timing or to supply timing to other subsystems. On the other hand, the conventional card is relegated to operate with its own ±100 ppm internal clock. Note that the conventional card is still able to use the clock from an external subsystem (for example the CPU) for data transmission but data reception is not coupled to the transmitter clock and it is also uncoupled to other transmitters in the network. This last feature it the one that defines IEEE 802.3 Ethernet as an asynchronous technology.

Synchronous Ethernet ability to accept or give timing signal makes this technology suitable for hierarchical synchronization. Here, the key element is the EEC which enables Ethernet nodes to accept or supply synchronisation to other Ethernet or TDM equipments. Thanks to this property, Synchronous Ethernet becomes a new building block of the synchronization network.

10.2.1.Ethernet Synchronization Messaging Channel

In SDH, the SSM provides traceability of synchronization signals and it is therefore required to extend the SSM functionality to Synchronous Ethernet to achieve full inter operability with SDH equipment.

In SDH, the SSM message is carried in fixed locations within the SDH frame. However, in Ethernet there is no equivalent of a fixed frame. The mechanisms needed to transport the SSM over Synchronous Ethernet are defined by the ITU-T in G.8264 in cooperation with IEEE. More specifically, the ESMC, defined by the ITU-T is based on the Organization Specific Slow Protocol (OSSP), currently specified in IEEE 802.3ay.

The ITU-T G.8264 defines a background or *heart-beat* message to provide a continuous indication of the clock quality level. However, event type messages with a new SSM quality level are generated immediately.

The ESMC protocol is composed of the standard Ethernet header for a slow protocol, an ITU-T specific header, a flag field, and a type length value (TLV) structure. The SSM encoded within the TLV is a four-bit field whose meaning is described in ITU-T G.781.

10.2.2.Synchronous Ethernet for the 1000BASE-T Interface

Historically, the 1000BASE-T is the first Ethernet interface that makes use of advanced modulation and encoding technology to enable simultaneous full duplex transmission over four twisted pairs in Cat. 5 cables.

Another property that makes different the 1000BASE- T interface to other Ethernet interfaces such as 1000BASE-X or 100BASE-TX is synchronization. The 1000BASE-T modulation is not compatible with asynchronous operation. During the auto-negotiation process, 1000BASE-T peers decide which transmission end becomes the master and which is the slave. This basic synchronization mechanism is suitable to achieve synchronization in the 1000BASE-T link but is not a mechanism it can be used for global synchronization of the Ethernet network. In fact, the 1000BASE-T synchronization mechanism constitutes a limitation to the operation of Synchronous Ethernet. One link that is operating as a 1000BASE-T slave is unable to accept a timing signal from the EEC and thus is unable to accept an arbitrary timing source.

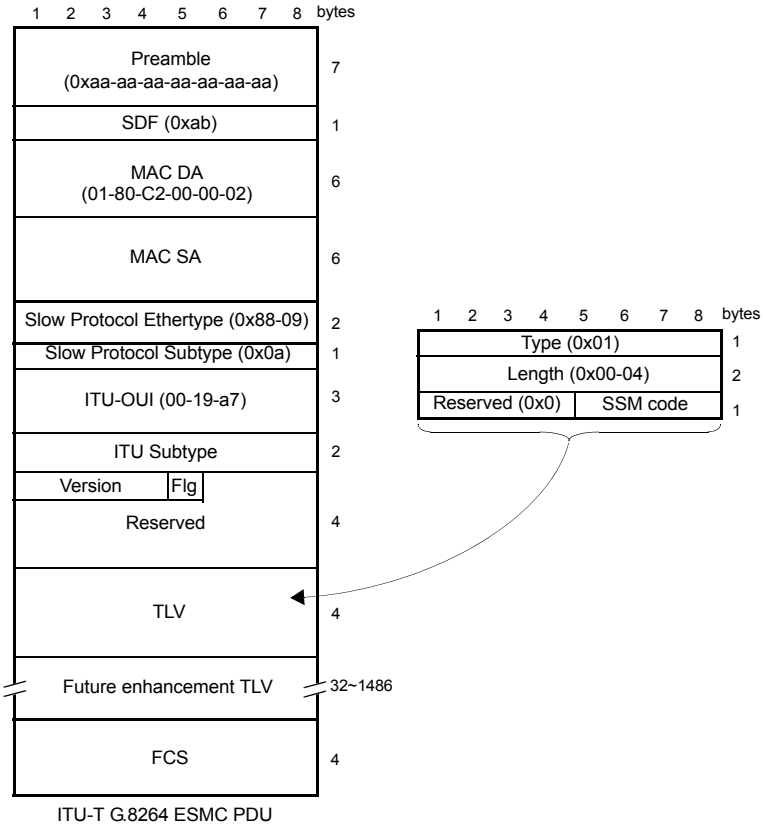

Figure 10.1: Ethernet Synchronization Message Channel (ESMC) protocol data unit.

This condition is different to other Ethernet interfaces where the transmission is synchronized to timing sources other than the reception clock and are therefore suitable to be added to a hierarchical synchronization network through the EEC without any special constrain.

Ethernet 1000BASE-T interfaces can still use Synchronous Ethernet but the EEC is always required to be connected to the master while the slave is constrained to operate with the timing signal recovered from the master. The result is that synchronization of 1000BASE-T is always unidirectional and propagating from the master to the slave

while in 1000BASE-X or 100BASE-T you can theoretically have one synchronization signal propagating in each transmission direction.

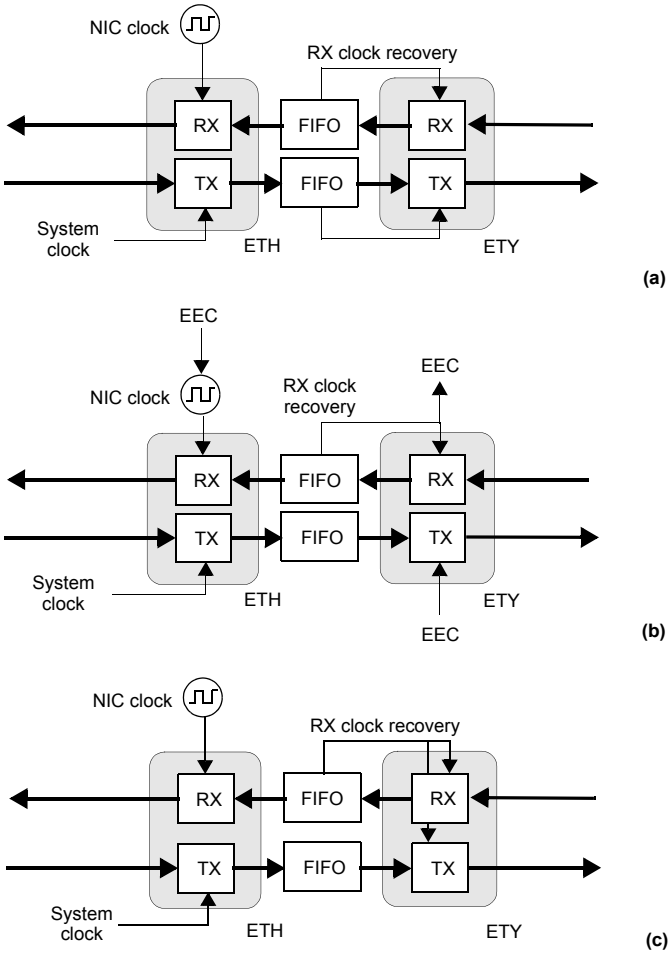

Figure 10.2: Ethernet synchronization models: (a) Conventional Ethernet. There are at least two clock domains. The transmitter and the receiver are in different clock domains (b) Synchronous Ethernet. All clocks are slaves of the EEC. The RX recovered clock could be used as an input for the EEC. (c) 1000BASE-T slave. The TX clock is always an slave of the RX clock.

10.3.Synchronous Ethernet Frequency Measurement

In conventional Ethernet, the received frequency is not a critical parameter as long as it is within the ±100 ppm range specified by the IEEE 802.3 standard. However, Synchronous Ethernet frequency is important because this frequency is used by external elements as a synchronization source. Ether.Genius and Ether.Sync include a frequency measurement that can be used to verify that the frequency offset of the received signal lies within acceptable limits. To enable the frequency measurement with Ether.Genius or Ether.Sync, follow these steps:

1. Make sure that your tester is connected to the network. The physical layer must be up and working in the correct test interface (See section 4.1.1).
   Note: if you are measuring frequency over the 1000BASE-T interface, force the slave role in the test port by disabling *100-FD* and *10-FD* and configuring *Clock role* to *Slave* (See section 2.2.4).
2. From the *Home* panel, go to *Results*,
   The port setup panel is displayed.
3. Select either *Port A or Port B* to enter in the port specific results menu.
4. Go to Frequency.
5. Check the *Frequency* and *Frequency deviation* results.

If you need increased accuracy for your frequency test, you can use an external clock for measuring. Different kinds of clock inputs are accepted by Ether.Genius and Ether.Giga, including TDM and GPS sources (See section 2.6)

10.4.Synchronous Ethernet Frequency Offset Generation

Ether.Genius and Ether.Sync can be used to impair the frequency of the test signal generated in test ports A and B with a frequency offset in the range of -125 ppm and +125 ppm. This test can be used to check how tolerant to frequency variations is a network element. Also, if you have a chain with various equipments transmitting a synchronization reference, you can replace the reference by the Ether.Genius / Ether.Giga test signal and test the ability of the chain to transmit frequency variations to the last element of the chain. The configuration procedure for frequency offset generation over Ethernet interfaces is as follows:

1. Make sure that your tester is connected to the network. The physical layer must be up and working in the correct test interface (See section 4.1.1).
   Note: if you are generating frequency offset over the 1000BASE-T interface, force the master role in the test port by disabling *100-FD* and *10-FD* and configuring *Clock role* to *Master* (See section 2.2.4).
1. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
2. Select *Port A* to enter in the port specific configuration.
3. Go to *Physical layer* to enter in the physical settings configuration panel.

4. Configure the *Frequency deviation (ppm)* with the frequency offset you want to generate between -125 and 125 ppm.

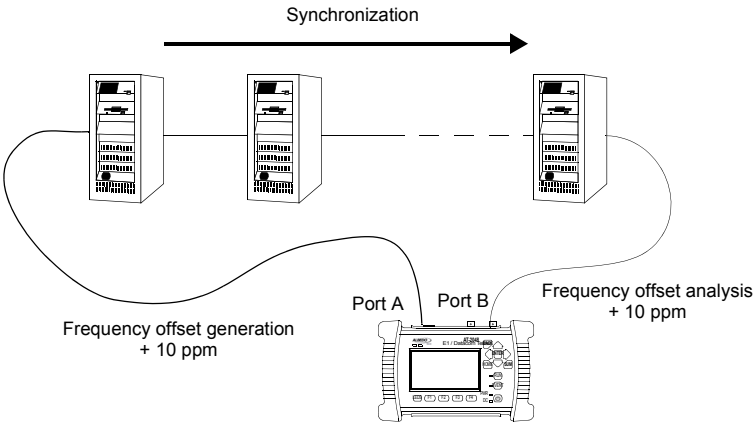

Figure 10.3: If Synchronous Ethernet is used for transmission of timing signals, network elements in the chain should be transparent to offset generation at is input and therefore any frequency offset should be recovered at the output of the chain.

10.5.Operation in IP Through Mode

You can use Ether.Genius or Ether.Sync to forward and impair the Ethernet frequency in pass-through mode. Specifically, Ether.Genius and Ether.Sync are capable of forwarding synchronization and traffic between port A and B and adding frequency offset to the signal forwarded to port A. The procedure to do that is as follows:

1. Make sure that your tester is connected to the network in pass-through mode. The physical layer must be up and working in ports A and B (See section 4.1.1).
   Note: if you are working with the 1000BASE-T interface, force the master role in the test A and slave role in test port B. Also, you have to disable the *100-FD* and *10-FD* auto-negotiation options for both ports (See section 2.2.4).
2. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
3. Configure *Mode* to *IP Through*.
4. From the *Home* panel, go to *Setup*,
   The port setup panel is displayed.
5. Go to *Reference clock*.
6. Configure *Input clock* to *Ethernet (Port B)*.

7. From the *Setup* panel, now go to *Port A* to enter in the port specific configuration.
8. Go to *Physical layer* to enter in the physical settings configuration panel.
9. Configure the *Frequency deviation (ppm)* with the frequency offset you want to generate between -125 and 125 ppm.

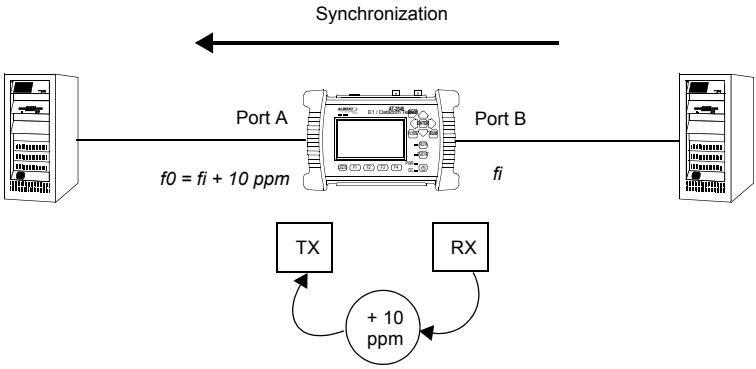

Figure 10.4: With the help of Ether.Genius or Ether.Sync it is possible to generate frequency offset in pass-through mode to stress the network.

Chapter 11
Test Management

This chapter describes all those features available in your test unit that are not directly related with configuring your tester or reading measurement results but they are important for proper test management. Specifically, configuration and result management, report generation and test platform settings are covered in the following sections.

11.1.Generating Reports

Users may want to generate reports based on their measurements. Reports are important to save results for later reference. Reports can be used to share a test result or to include results in documents.

Depending on the purpose of the report, users have different ways to generate and store them. Ether.Genius / Ether.Sync / Ether.Giga offer maximum flexibility and at the same time simplicity when configuring reports. Follow these steps to configure and generate a report:

1. From the *Home* panel, go to *File*,
   The tester file manager base menu is displayed.
2. Select *Report files* to go to the report file settings
3. Enable report generation by means the *Generate reports* control.
4. Set the *Report format*, *Report named after* and *Report header* fields.
5. If you have set *Report named after* to *User ref.+sequence*, configure the *User reference* field to the desired sequence.
6. Optionally, if you have configured *Report named after* to *User ref.+sequence* or *Serial no.+sequence*, enter the *Next sequence number* to be applied to the next report.
7. Set the correct action to be carried out when the internal storage is full.

If report generation is enabled, a new report is generated each time a test finishes either by pressing the run button or automatically. Reports are available as standard

text or PDF files from the USB slave connector and they can be exported through the USB master port, the SD card reader or the web interface.

Table 11.1: System Settings Panel

Setting	Description
Internal memory	Displays report files stored in the internal tester memory. Ether.Genius / Ether.Sync / Ether.Giga can store up to 50 report files.
External devices	Displays report files stored in external devices (SD memory card, USB memories or drives) connected to the tester. The amount of files stored in an external device is only limited by the device capacity.
Generate reports	Enables / Disables report generation.
Report format	Selects the report format for future reports. • *PDF*: Reports are generated using the portable document format (PDF). Use this configuration if you want to make difficult for anyone to modify the report. • *Plain text*: Reports are text documents which can be edited with any text editor. Use this configuration if you want to modify the report or include it in a wider document.
Report named after	This control enables the user to choose between different templates for the report name. There are three different templates to choose: • *Start time*: The report is identified by a time stamp that contains both data and time with the following format: yyyy-MM-dd-hhmmss. • *User ref. + sequence*: The report name is set to a user configurable string plus a sequence number that is incremented for each new test. • *Serial no. + sequence*: The report name is set to the tester serial number plus a sequence number that is incremented for each new test.
User reference	This could be any alphanumeric string containing upper case letters, lower case letters and numeric digits. This field makes sense only if the report name format is *User ref. + sequence*.
Next sequence	Displays and configures the sequence number that will be assigned to the next report to be generated. This field makes sense only if the report name format is *User ref. + sequence* or *Serial no. + sequence*.

Table 11.1: System Settings Panel

Setting	Description
Report header	This menu item enables you to configure report data that will be stored with the test result. These data identify the report, customer, and also includes some other relevant information. • *Customer*: Field that can be used to set the company where the test report applies. • *Department*: This field can be used to identify the department where the user that has carried out the tester. belongs. • *Company*: This is the field that identifies the company that carries out the test. • *Location*: This describes where the test results from the network were recorded. • *Operator*: This field may contain the name of the operator that owns the network infrastructure where the test was run.
Maximum reports	Displays the maximum number of reports that can be stored within the tester internal memory. Currently this number is limited to 50 files.
Action when disk full	Action to be carried out when the maximum reports limit is reached. There are three possible choices here: • *Block measurements*: No new measurements can be run when the internal memory is full • *Stop report generation*: New measurements are run even if the internal memory is full but no reports are generated for them. • *Delete oldest reports*: When the maximum available capacity is reached, new files replace the older ones. Use this action with care. No warning is displayed when old reports are deleted.

11.2.File Management

Ether.Genius / Ether.Sync / Ether.Giga stores configurations and reports in files. These files can be deleted, renamed or exported to an external USB memory or SD card. Configurations can be shared between different Ether.Genius / Ether.Sync / Ether.Giga units by means compatible storage devices. Report files can be included to documents, sent by e-mail or printed.

11.2.1.Saving Configurations

To store the current configuration follow these steps:

1. From the *Home* panel, go to *File*,
   The tester file manager base menu is displayed.
2. Select *Configuration files* to go to the configuration file settings.
3. Select the location to save the configuration: *Internal memory*, or *External devices*.
   Note: If you select E*xternal devices*, you will be asked to choose the specific storage device (USB device or SD card).
   Note: If there is no external device connected to the Ether.Genius / Ether.Sync / Ether.Giga unit, a *No devices present* popup panel is displayed.
4. Press the *Save* (F2) contextual button.
5. Enter a file name for the configuration file that is going to be saved and confirm with the *Done* (F4) contextual button.

11.2.2.Renaming Files

Both configuration and report files can be renamed after they are created. To rename files follow these sequence:

1. From the *Home* panel, go to *File*,
   The tester file manager base menu is displayed.
2. Select *Configuration files* or *Report files*.
3. Select the location of the file you want to rename: *Internal memory*, or *External devices*.
   Note: If you select E*xternal devices*, you will be asked to choose the specific storage device (USB device or SD card).
   Note: If there is no external device connected to the Ether.Genius / Ether.Sync / Ether.Giga unit, a *No devices present* popup panel is displayed.
4. Select the file you want to rename with the help of the cursors and the ENTER button.
   Note: You can select several files in the list, but renaming of many files at the same time is not allowed.
5. Press the *Rename* contextual button.
6. Enter the new file name for the selected configuration or report file with the alphanumeric keyboard. Confirm with the *Done* (F4) contextual button.

11.2.3.Deleting Files

With the file manager you can delete files that are not needed anymore. To do that follow these steps:

1. From the *Home* panel, go to *File*,
   The tester file manager base menu is displayed.
2. Select *Configuration files* or *Report files*.

3. Select the location of the file you want to delete: *Internal memory*, or *External devices*.
   Note: If you select E*xternal devices*, you will be asked to choose the specific storage device (USB device or SD card).
   Note: If there is no external device connected to the Ether.Genius / Ether.Sync / Ether.Giga unit, a *No devices present* popup panel is displayed.
4. Select the file you want to delete with the help of the cursors and the ENTER button.
   Note: You can select several files in the list at the same time.
5. Press the *Delete* contextual button.
6. Enter the new file name for the selected configuration or report file with the alphanumeric keyboard. Confirm with the *Done* (F4) contextual button.

11.2.4.Exporting Files to External Devices

Configuration and report files can be exported to external devices like USB memories or SD cards. The procedure is as follows:

1. From the *Home* panel, go to *File*,
   The tester file manager base menu is displayed.
2. Select *Configuration files* or *Report files*.
3. Select *Internal memory*, to list the files currently stored in the Ether.Genius / Ether.Sync / Ether.Giga unit.
4. Select the files you want to export with the help of the cursors and the ENTER button.
5. Press the *Export* contextual button.
   A popup menu to select the external device where the files will be exported is opened.
   Note: If there is no external device connected to Ether.Genius / Ether.Sync / Ether.Giga, a *No devices present* popup panel is displayed.
6. Select an external device, confirm, and wait for the files to be copied.
7. Remove the USB storage device or SD card from the unit.

11.2.5.Importing Configurations

If you have a configuration file from a compatible tester you can import and load this file in your unit to reproduce similar measurements. This is the procedure you have to follow:

1. From the *Home* panel, go to *File*,
   The tester file manager base menu is displayed.
2. Select *Configuration files* to go to the configuration file settings.
3. Select *External devices* to list the files currently stored in the external device.
   A popup menu to select the source external device is opened.

Note: If there is no external device connected to the Ether.Genius / Ether.Sync / Ether.Giga unit, a *No devices present* popup panel is displayed.

4. Select the configuration files you want to import with the help of the cursors and the ENTER button.
5. Press the *Import* contextual button, confirm, and wait for the files to be copied from the internal memory.
6. Remove the USB storage device or SD card from the unit.

11.2.6.Using the Embedded Web Server

As an alternative of using a USB external storage device or an SD card for file management, Ether.Genius / Ether.Sync / Ether.Giga has a web interface that can be used for the same purpose.

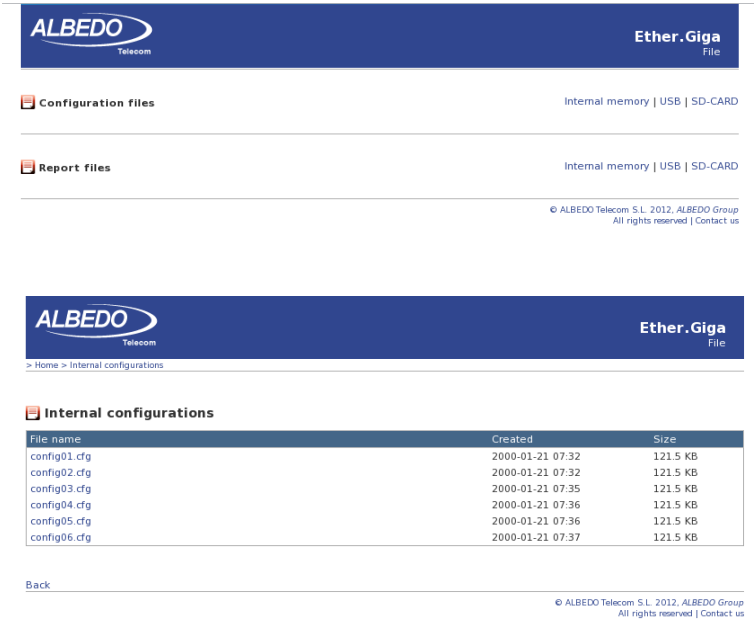

Figure 11.1: Ether.Genius / Ether.Sync / Ether.Giga web interface: (a) Home panel (b) Configuration management panel.

The web interface can be used for downloading configurations and reports from a remote computer without using any accessory other than an standard network connection. Currently, the web interface does not support file uploading but for this purpose, the USB and SD interfaces are still available.

To use the web interface you need to connect the platform network connector to the management network and configure the management Ethernet interface (See section 11.4.1). Once you have done this, follow this procedure:

1. Open a browser in a computer with network connection.
2. Type the IP address you have assigned to the tester in the browser destination URL.
   The web interface home panel is displayed in the Internet browser.
3. Choose the files you want to display (*Configuration flies*, *Report files* or any other if available) and the location of these files (*Internal memory*, *USB*, *SD-CARD*) and press to the correct hyper link.
   A list with the available files for the selected category is displayed in the web browser.
4. Select the file you want to download it to the local computer.
   The web browser displays a dialogue that requests your configuration to download the selected file. If you accept, the file will be downloaded.

11.3.Programming Tests

Ether.Genius / Ether.Sync / Ether.Giga is able to start and finish tests without direct user intervention. All automatic testing features are included within the *Autostart/stop* menu

Follow these steps to program an automatic measurement in the test unit.

1. From the *Home* panel, go to *Test*,
   The test configuration panel is displayed.
2. Select *Autostart/stop* to enter in the automatic test programming menu.
3. If you want the automatic test to start at a specific date and time set *Start mode* to *Auto* and enter the start date and time in *Start time*.
   Note: Manual start has precedence over autostart. That means that if a tester is started by pressing RUN but there is an automatic test programmed the manual test will start anyway.
4. If you want the automatic test to stop at a specific time after it has started set *Stop mode* to *Auto* and enter test duration with the help of the *Duration* and *User duration* controls.
   Note: Manual stop has precedence over autostop. That means that if a tester is

stopped by pressing RUN but there is an automatic test programmed the manual test will stop anyway.

Table 11.2: System Settings Panel

Setting	Description
Start mode	Configures the start test mode. There are two different choices here: • *Manual*: The test starts when there is not an ongoing test and the RUN key is pressed. • *Auto*: The test starts at a configured date and time without the need of pressing any key.
Start time	Enter the start date and time for the next automatic measurement with the following format: *dd/MM/yyyy hh:mm:ss*. To configure Start time, you have to set *Start* mode to *Auto* before.
Stop mode	Configures the stop test mode. There are two possibilities: • *Manual*: The test finishes when there is an ongoing test and the RUN key is pressed. • *Auto*: The test finishes when a configurable test duration is reached. This mode does not require user intervention once the duration has been set and the measurement has started.
Duration	Sets the duration of the next measurement. The available test durations are: 15 minutes, 1 hour, 1 day, 7 days, 30 days or user configurable duration. Setting up *Duration* requires previous configuration of *Stop Mode* to *Auto*.
User duration	Sets the duration of the next measurement when *Stop mode* has been configured to *Auto* and *Duration* to *User*. The duration has to be entered in a *hh:mm:ss* format.
Last started on	Displays the date and time when the last measurement was started.
Last stopped on	Displays the date and time when the last measurement was stopped. If there is an ongoing test, the value of this field is empty.
Last power down on	Displays the date and time when the tester was powered down for last time.

11.4.Using the System Menu

The System menu includes platform wide settings organized in four different submenus:

- *General settings*: This menu includes controls to manage the way the user interface behaves and how the information is presented.
- *Network configuration*: Includes the IP configuration corresponding with the platform NIC.
- *System information*: This menu has the test unit model name and serial number and software, firmware and hardware versions.
- *Licensed options*: This is a menu that displays the software versions installed in the tester and enables their management.

Table 11.3: System Settings Panel

Setting	Description
Brightness (%)	Sets the screen brightness from 10% to 100%. Within the *Brightness* panel, the left and right cursors are used to set the correct value and a contextual key (*Done*) is used to confirm selection.
Keyclick	Enables or disables the keyclick. The keyclick is a sound that is played each time a key is pressed.
Language	Selects the user interface language. Menus, selection lists and results are presented in the language selected here. The languages currently available are English and Spanish.
Clock setup	Configures the system time and date. You can either type the correct date and time manually or let the equipment to retrieve the correct values from a *Network Time Protocol* (NTP) server.
Time display	Select the way the time is displayed in the graphical user interface. One of the following has to be selected: • *Elapsed*: Time from the beginning of the test is displayed with the following format *hh:mm:ss*. If there is not an ongoing test, then the duration of the last test is shown • *Absolute*: The current date and time is displayed with the following format: *dd/MM/yyyy hh:mm:ss*.
Screensaver	Sets or unsets the screensaver. The screensaver reduces power consumption and increases operation time under battery operation.

Table 11.3: System Settings Panel

Setting	Description
Screensaver delay	Configures the delay to switch the screensaver on. The backlight brightness is set to a low value once the time configured here has finished. The display backlight is switched off after twice the screensaver delay. The available configuration values for this item are: 10s, 30s, 1min, 2min, 5min, 10min, 20min.
Remote control	Enables or disables the Ethernet / IP remote control. The remote control is an optional feature that enables remote users to use the tester from a computer running VNC.
Remote control password	Configures a password for the remote control. Any alphanumeric string should be accepted. Use the same password in the remote VNC client to access to the tester user interface.

This section supplies a description of the *General settings* menu and *System information* menu. To learn how to configure and use the network interface or how to install licenses for new software options go to the sections specifically dedicated to these topics.

Table 11.4: System information panel

Setting	Description
Model Name	Shows the test unit model name: Ether.Genius, Ether.Sync or Ether.Giga.
Serial number	Displays the test unit serial number. It is a 8 character alphanumeric string
Software release	Displays the current software release.
Hardware release	Displays the current hardware release.
Firmware release	Displays the current firmware release.
PM release	Displays the current power management release.

11.4.1.Using the Network

The platform network interface is currently user for three different purposes:

- The *Ethernet / IP remote control*. This feature enables any user to access to the equipment form a remote location, configure a test, run it and display the results.
- The *Web interface*: This is used to retrieve reports configurations or any other file available in the tester internal memory or attached storage device.

- *Maintenance and factory configuration*: The ALBEDO Telecom staff use the Ethernet interface to configure or verify the equipment in the factory. This feature is not available to ordinary users.

Table 11.5: Network Configuration Panel

Setting	Description
Ethernet interface	Configuration menu for the platform network interface. This menu can be used to configure the interface IP address and mask either automatically (DHCP) or statically.
Wireless interface	Configuration menu for the platform wireless network interface. This menu is used to set the radio parameters for the interface such as the SSID and the network parameters like the IP address and mask. The wireless interface requires a compatible WiFi adapter for the USB port. This adapter is supplied by ALBEDO as an optional accessory.
Gateway address	IP address corresponding to the IP default gateway in four dotted format. There is only one default gateway for all the network interfaces (wired and wireless). By setting up this field, the user decides which management port is used by the system to reach remote networks. It is possible to configure the gateway address automatically if either the wired or the wireless interfaces are configured to get an IP profile through the DHCP protocol.
DNS address	DNS server address used by the platform management ports to resolve domain names. It is possible to configure the DNS address automatically if either the wired or the wireless interfaces are configured to get an IP profile through the DHCP protocol.

To configure and use the Ethernet platform interface follow these steps:

1. From the *Home* panel, go to *System*,
   The general system menu is displayed in the screen.
2. Select *Network configuration* to display the network configuration and management menu.
3. Go to the *Ethernet interface*.
4. Enable the platform network interface with the *Enable interface* control.
5. Enable DHCP with the *Use DHCP* control if you want to let DHCP to configure your IP settings automatically or disable it to configure an static IP profile.

6. If you are not using DHCP, enter correct values for the *Static IP address* and *Static network mask*.
7. Leave the *Ethernet interface* panel with the ESC key.
8. If you are not using DHCP, configure the *Gateway address* and *DNS address*.
9. Connect the platform Ethernet connector (platform panel, RJ-45 connector with the *Ethernet* label) to the management network.
10. Optionally, check from a remote computer that the equipment is responding to ping requests.

Table 11.6: Ethernet Interface Configuration

Setting	Description
Enable interface	Enables or disables the network interface. Note that the link led placed in the Ethernet platform connector is lit even if the interface is not enabled.
Use DHCP	Configures the mechanism used to set the interface IP address and mask (and also other system-wide settings like the gateway address and the DNS server). If *Use DHCP* is enabled, the IP profile is configured automatically using a DHCP server installed in the network. Otherwise, the user has to enter the IP address, mask, default gateway and DNS address by hand.
Static IP address	Static IP address assigned to the interface in a decimal four dotted format. This setting makes sense only if *Use DHCP* is not enabled.
Static network mask	Static network mask in a decimal four dotted format. This setting makes sense only if *Use DHCP* is not enabled.
Leased IP address	Current DHCP-assigned IP address in a decimal four dotted format. This is a read-only field that cannot be directly configured by users This setting makes sense only if *Use DHCP* is enabled.
Leased network mask	Current DHCP-assigned network mask in a decimal four dotted format. This is a read-only field that cannot be directly configured by users This setting makes sense only if *Use DHCP* is enabled.
Ethernet address	48-bit physical address of the NIC attached to the test unit. This address is assigned to the NIC when it is manufactured and it cannot be changed later.

11.4.2.Installing Software Options

New software for Ether.Genius / Ether.Sync / Ether.Giga can be licensed after the unit as been purchased when new testing needs arise. To install new software options for your unit follow this procedure.

Table 11.7: Licensing

Setting	Description
Licensed options	Shows a list with all the software options currently available in your test unit.
License number	8-digit hexadecimal number provided by ALBEDO Telecom that identifies the software options to be added to your unit. Enter your license number in this field before adding the new software options to your test unit.
License key	8-digit hexadecimal number provided by ALBEDO Telecom that enables secure management of the software options installed in your test unit. Enter the license key in this field before adding the new software options to your test unit.
Activate	Set this field to Yes to add new software options to your tester. You have to enter the *License number* and the *License key* before adding new options.
Status	Displays the result of the software option activation operation performed by enabling the *Activate* field.

1. Contact with your local sales representative to purchase software options for your test units.
   You will receive one license number and one license key for each tester you want to upgrade.
2. From the *Home* panel, go to *System*,
   The system configuration panel is displayed.
3. Select *Licensing* to enter in the software upgrade menu.
4. Enter the number and key supplied by your ALBEDO Telecom representative in *License number* and *License key*.
5. Enable the new software options with the *Activate* control.
6. Check that the upgrade has been successful with the help of the *Status* control.

11.4.3.Using NTP for System Clock Synchronization

The system clock controls the date and file assigned to configuration and report files and controls the autostart / stop function. Ether.Genius / Ether.Sync / Ether.Giga users

can manually set the system time and date by entering the correct values but they can also synchronize the clock with an external NTP server. The NTP server must be available through the platform Ethernet port.

Table 11.8: Time Source Configuration Options

Setting	Description
Date	This is used to configure the current date. The date is used for the *Autostart/stop* features and other purposes. The date has to be entered with the following format: *dd/MM/yyyy*. You can only modify the date if the current time source has been set to *Manual*.
Time	This is used to configure the current time. The time is used for the *Autostart/stop* features and other purposes. The time has to be entered with the following format: *hh:mm:ss*. You can only modify the time if the current time source has been set to *Manual*.
Time source	It is either *Manua*l or *NTP.* The meaning of each configuration option is as follows: • *Manual*: The user configures the *Date* and *Time* fields manually. System date and time is controlled by the internal clock. • *NTP*: An external *Network Time Protocol* (NTP) server controls the value of the *Date* and *Time* fields. The system is synchronized with the server each time the equipment is restarted or when a new capture is run.
UTC offset (hours)	This is the time difference in hours between your local time zone and the *Universal Time Coordinated* (UTC) time zone This setting makes sense only if *Time Source* has ben configured to *NTP*.
UTC offset (min)	This is the value to be configured when the time offset between your local time zone and the UTC zone is not an integer value of hours. This setting makes sense only if *Time Source* has ben configured to *NTP*.

Table 11.8: Time Source Configuration Options

Setting	Description
NTP server	IP address or domain name corresponding to the NTP server you want to use to synchronize your unit. If you want to use a domain name for the server you need to make sure you have configured a DNS server in your network settings (See section 11.4.1). This setting makes sense only if *Time Source* has ben configured to *NTP*.
NTP status	Displays the current status of the NTP server configured in the *NTP server* field. It is one of the following: • Server not available: The server configured in NTP server is not available. Make sure that your network interface is properly configured and that the server is accessible. • *Synchronized*: The equipment is correctly synchronized with the external server configured in *NTP server*.. • *Waiting*: The test unit is still waiting for a reply from the remote server. This is a not editable field. It is only active if *Time Source* has ben configured to *NTP*.

To configure the system time and date in your Ether.Genius / Ether.Sync / Ether.Giga unit follow these steps:

1. From the *Home* panel, go to *System*,
   The general system menu is displayed in the screen.
2. Select *General settings* to display the platform-wide configuration.
3. Go to *Clock setup*
4. Configure *Manual* or *NTP* time and date with the help of the *Time source* field.
5. If you have configured *Time source* to *Manual* in the previous step, configure the *Date* and *Time* field to your local time and date. If *Time source* has been configured to *NTP*, enter the *UTC offset (hours)*, *UTC offset (min), NTP server* and wait for the equipment to establish synchronization with the server.

11.5.Using the Remote Control

The remote control application constitutes a remote graphical user interface that reproduces pixel by pixel the tester screen in virtually any remote device supporting the VNC protocol. This includes not only computers but also smartphones or tablets. The only requirements for the controlling devices are:

- IP connectivity with the tester. Any IP connection including Ethernet, WiFi and 3G should work.
- They must have a VNC client installed. Currently, there are VNC clients for most OS in the market. Some of them are free.

The remote control is an optional feature for Ether.Genius / Ether.Sync / Ether.Giga that is supplied by ALBEDO Telecom with an special license.

Before using the remote control you need to configure the platform Ethernet interface and connect the equipment to the management network (See section 11.4.1). Once this is done, follow this procedure to use the remote control:

1. From the *Home* panel, go to *System*,
   The system configuration panel is displayed.
2. Select *General settings* to display miscellaneous system-wide settings, including the ones referred to the remote control.
3. Enable the remote control with the help of *Remote control*.
4. Optionally, supply a password with *Remote control password*. The password you configure here will be requested in all incoming VNC connections.
5. In the controlling device, run the VNC client and enter the password you have configured in *Remote control password* if you are requested to do so.
6. Use the keyboard (navigation through the mouse is not available in the remote control) to browse the instrument panels, start measurements, insert events or any other action.

Table 11.9: Remote Control Keys

Key	Description
Up, Down, Left, Right	These keys are equivalent to the cursor keys in the tester local interface. They move the focus through the different fields available in the current panel and they also help with the navigation through different panels.
Home	It is equivalent to the HOME key. It displays the *Home* panel.
Esc	It is equivalent to the Esc key. It leaves the current panel and displays the previous one in the panel hierarchy.
Enter	It is equivalent to the ENTER key. It confirms settings.
Ctrl+L	It is equivalent to LEDS. It displays the *Leds* panel.
Ctrl+S	It is equivalent to SUM. it displays the *Summary* screen
Ctrl+R	It is equivalent to RUN. It starts / stops a measurement
Ctrl+E	It is equivalent to EVENT. It starts / stops event insertion
F1, F2, F3, F4	They are equivalent to the F1, F2, F3, F4 contextual keys. The purpose of these keys depend on the current screen.

Appendix A
Technical Specification

A.1. General

- Operation over two Gigabit Ethernet physical interfaces based either on SFPs or RJ45 connectors.
- Traffic generation and analysis features up to 1 Gb/s (1.5 millions of frames, if frame size is set to 64 bytes).

A.2. Operation Modes

- *L1 Endpoint operation: The equipment generates PCS codes and L1 BER measurement patterns.*
- *Ethernet Endpoint operation*: The equipment generates and receives Ethernet PCS codes and Ethernet frames.
- *IP Endpoint operation*: The equipment generates and receives IPv4 datagrams.
- *Through operation*: Traffic is forwarded between port A and B.

A.3. Ethernet PHY

- Supported interfaces (SFP): *10BASE-T*, *100BASE-TX*, *100BASE-FX*, *1000BASE-T*, *1000BASE-SX*, *1000BASE-LX*, *1000BASE-ZX*.
- Supported interfaces (RJ-45 ports): *10BASE-T*, *100BASE-TX*, *1000BASE-T*.
- On/off laser control for optical interfaces.

A.3.1. Auto-Negotiation

- Negotiation of bit rate. Allow 10 Mb/s, allow 100 Mb/s, allow 1000 Mb/s.
- Negotiation of *Master* and *Slave* roles in the 1000BASE-T interface.
- Ability to disable auto-negotiation and force line settings.

A.3.2. Power over Ethernet

- PoE (IEEE 802.3af-2003) and PoE+ (IEEE 802.3at-2009) detection.
- PoE interfaces: 10BASE-T, 100BASE-T and 1000BASE-TX through attached RJ-45 ports A and B.
- PoE pass-through when the equipment is configured in transparent (through) operation mode.

A.4. Synchronous Ethernet

- Supported interfaces: 100BASE-TX and 1000BASE-T through the attached RJ-45 ports. 1000BASE-SX, 1000BASE-LX, 1000BASE-ZX and 1000BASE-BX through external SFP.

A.4.1. Operation

- Analysis of synchronous Ethernet signal in Ethernet endpoint, IP Endpoint and Through modes, generation of synchronous Ethernet signal in Ethernet endpoint and IP Endpoint modes. Transparent synchronous Ethernet pass-through in Through mode.
- Configuration of internal, external or recovered clock in Ethernet interfaces.
- Fixed freq. offset generation on transmitted signals with maximum value of ±125 ppm (resolution 0.001 ppm) as per ITU-T O.174 (11/2009) 8.2.1.

A.4.2. Analysis

- Measurement of the line frequency (MHz), frequency offset (ppm) and frequency drift (ppm/s) as specified in ITU-T O.174 (11/2009) clause 10.

A.5. Clock References

- Internal time reference better than ±2.0 ppm. Optional internal reference better than ±0.2 ppm.
- Ethernet input through Port A or Port B over any valid electrical / optical synchronous Ethernet interface.
- 10 MHz, 2048 Mb/s, 2048 MHz, 1544 Mb/s, 1544 MHz input through Port C (balanced or unbalanced) or through DTE port (balanced interface, RJ-48 connector adapter).
- One-pulse-per-second (1 pps) synchronization input through DTE port (balanced interface, RJ-48 connector adapter).

- 2048 kHz reference output through Port C (balanced or unbalanced), 1 pps through DTE port (balanced interface, RJ-48 connector adapter).

Table 1: Clock reference input levels

Type	Pulse	Mode	Level (Max)	Level (Min)	Duty cycle
1544, 2048 kHz	Unipolar, square, senoidal	Endpoint	5 Vpp	500 mVpp	40 - 60 %
		Monitor	0.5 Vpp	100 mVpp	40 - 60 %
1544, 2048 kb/s	AMI, HDB3	Endpoint	Nom. G.703	-12 dB (line)	40 - 60 %
		Monitor	Monitor -20 dB		40 - 60 %
10 MHz	Unipolar, square, senoidal	Endpoint	5 Vpp	1 Vpp	45 - 55 %
		Monitor	1 Vpp	100 mVpp	45 - 55 %

A.6. Ethernet MAC

- Frame formats: *DIX*, *IEEE 802.1Q*, *IEEE 802.1ad*.
- Support for Jumbo frames with MTU up to 10 kB.
- Setting of *source* and *destination MAC addresses*. Destination addresses can be configured as a single value or as a range.
- Setting of the *Type / Length* value.
- Configuration of the *VID* and *priority codepoint* in VLAN modes.
- In Q-in-Q / IEEE 802.1ad modes, configuration of the S-VLAN *VID*, *DEI* and *priority codepoint*. Configuration of the C-VLAN *VID* and priority *codepoint*.
- Configuration of the *frame size*.

A.7. MPLS

- MPLS generation and analysis in *IP Endpoint* mode. MPLS analysis in *IP through* mode.
- Support of a single and double label stack (*Top* and *Bottom* labels). Label formatting follows RFC 3032.
- Configuration of the TTL, traffic class and label value for *Top* and *Bottom* MPLS headers.

A.8. IPv4

- Configuration of *source* and *destination IPv4 addresses*. Destination addresses can be configured as a single value or as a range.

- Configuration of *DSCP* CoS labels, *TTL* and *transport protocol*.
- If transport protocol is UDP, support of UDP frame with *source* and *destination port* configuration.

A.9. Traffic Generator

- Generation over 8 independent streams. Each stream has its own specific bandwidth profile and payload / pattern configuration.

A.9.1. Bandwidth Profiles

- Generation modes: *Continuous*, *Periodic burst*, *Ramp* and *Random*.

A.9.2. Test Patterns and Payloads

- Layer 2-4 BER test patterns: *PRBS* 2^{11}*-1*, *PRBS* 2^{15}*-1*, *PRBS* 2^{20}*-1*, *PRBS* 2^{23}*-1*, *PRBS* 2^{31}*-1* along with their inverted versions and user (32 bits). These patters apply to stream 1 only.
- Test payload for SLA tests.
- *All zeros* test pattern.
- NCITS TR-25-1999 RPAT, JPAT and RPAT for L1 BER tests.
- IEEE 802.3, Annex 36A HFPAT, LFPAT, MFPAT, LCRPAT, SCRPAT for L1 BER tests.

A.10.Event Insertion

- Insertion of TSE, FCS errors, Undersized frames and IPv4 checksum errors.
- Insertion modes: Single, burst, rate and random.

A.11.Filter

- Up to 8 simultaneous filters can be applied to the traffic.
- The equipment supports a generic filter which can select frames by using a *16 bit mask* and an arbitrary *offset* defined by the user.

A.11.1. Ethernet Selection

- By *source* and *destination MAC addresses*. Selection of MAC address sets with masks.
- By *Type / Length* value with selection mask.
- By *C-VID* and *S-VID* with selection mask.
- By *service* and *customer priority codepoint* value with selection mask.

A.11.2. MPLS Selection

- Separated filters to account for the top and bottom MPLS headers.
- Filtering by label value.
- Filtering by traffic class.

A.11.3. IPv4 Selection

- Selection by I*Pv4 source* or *destination* address. It is possible to select address sets by using masks.
- Selection by *protocol*.
- Selection by *DSCP value*.

A.11.4. IPv6 Selection

- Selection by IPv6 source or destination address (or both at the same time). It is possible to select address sets by using masks.
- Selection by IPv6 flow label.
- Selection based on the next header field value.
- Selection by DSCP value.

A.11.5. UDP Selection

- Selection by *UDP port*. Either as a single value or a ranges.

A.12. PHY Results

A.12.1. Cable Tests

- *Optical power* measurement (transmitted and received power) over compatible SFP transceivers.
- Inactive links: *Open/short fault indication* and *distance to fault* in metres (accuracy: 1 m, range 100 m).
- Active links: current local port *MDI / MDIX* status, .cable wiring (straight, crossed), polarity (positive, negative) pair skew (1000BASE-T only), crosstalk.

A.12.2. Auto-Negotiation

- *Bit rate* and *duplex mode*.

A.12.3. SFP

- *SFP presence*, current *interface*, *vendor*, and *part number*.

A.12.4. Power over Ethernet

- Type of PoE: PoE (IEEE 802.3af), PoE+ (IEEE 802.3at), none.
- PoE voltage between pairs 1-2 / 3-6 and 4-5 / 7-8 in endpoint test. Voltage and current in pairs 1-2 / 3-6 and 4-5 / 7-8 in through mode.

A.13.Frame Analysis

- Support of *local one-way* (port A-port B) and *two-way* (port A-port A) test modes.
- Separate traffic statistics for Port A and B.

A.13.1. Ethernet Statistics

- Frame counts: *Ethernet*, *VLAN*, *IEEE 802.1ad frames*, *Q-in-Q frames*, *control frames*, *pause frames*.
- Frame counts: *unicast*, *multicast* and *broadcast*.
- Basic error analysis: *FCS errors*, *undersized frames*, *oversized frames*, *jabbers*.
- Frame size counts: *64 or less*, *65-127*, *128-255*, *256-511*, *512-1023*, *1024-1518*, *1519-1522*, *1523-1526* and *1527-MTU* bytes.

A.13.2. MPLS Statistics

- MPLS stack length: minimum, maximum.

A.13.3. IP Statistics

- Packet counts: *IPv4 packets, IPv6 packets*.
- Packet counts: *unicast*, *multicast* and *broadcast*.
- *UDP packets*, *ICMP packets*.
- *IPv4 errors*, *IPv6 errors.*
- *UDP* errors.

A.13.4. Bandwidth Statistics

- Ethernet traffic statistics expressed in bits per second, frames per second and a percentage of the nominal channel capacity.
- *Unicast*, *multicast* and *broadcast* traffic figures expressed in frames per second units.
- IPv4 and IPv6 statistics (frames per second, bits per second and percentage).
- UDP traffic (frames per second, bits per second and percentage).

A.13.5. SLA Statistics

- Multistream SLA analysis.
- Delay statistics: ITU-T Y.1563 *FTD* (current, minimum, maximum, and mean values).
- Delay variation statistics: ITU-T Y.1563 *FTD* (standard deviation), ITU-T Y.1563 *FDV* (peak), RFC1889 / RFC 3393 *jitter* (current, maximum and mean values).
- Frame loss: ITU-T Y.1563 *FLR*.
- Duplicated packets, out -of -order packets (RFC 5236).
- Availability statistics: SES and ITU-T Y.1563 *PEU*.

A.13.6. BER

- *Bit error count*, *seconds with errors*, *bit error ratio* (BER).
- *Pattern losses*, *pattern loss seconds*.

A.13.7. Network Exploration

- *Top talkers* statistics: Displays the 16 most common source MAC / IPv4 / IPv6 addresses.
- *Top VID* (IEEE 802.1Q) or *C-VID* (IEEE 802.1ad): Displays the 25 most common VID / C-VID tags.

A.14. Automatic Tests

- The equipment supports automatic normalized tests defined in IETF RFC 2544 and ITU-T Y.1564 (eSAM).
- Support of *local one-way* (port A - port B) and *two-way* (port A - port A) tests.
- Support of Ethernet and IP test modes.

A.14.1. IETF RFC 2544 Test

- Support of RFC-2544 *throughput*, *frame-loss*, *latency*, *back-to-back* and *recovery time* tests.

A.14.2. eSAM Test

- Testing of up to eight services (non-colour aware mode) or up to four services (colour aware mode).
- Configuration of the CIR and *EIR* for each service.
- Configuration tests (CIR, EIR and policing) with *FTD*, *FDV*, *FLR* results for each service.
- Performance test with *FTD*, *FDV*, *FLR* and *availability* results for all services.

A.15. Port Loopback

- Layer 1-4 loopback.
- Loop frames matching current filtering conditions or loop all frames in layer 2-4 loopbacks.
- Loop controls for broadcast and ICMP frames.

A.16. Ping and Trace-route

- Generation of on demand *ICMP echo request* (RFC 792) messages with custom destination IP address, packet length and packet generation interval.

- Analysis of *ICMP echo reply* (RFC 792) messages with measurement of round trip time and lost packets.
- Analysis of *ICMP Time-To-Live Exceeded* and *ICMP Port unreachable* replies received in the trace-route test.

A.17.PTP / IEEE 1588

- Operation: IEEE 1588-2008 transparent, non-intrusive monitoring in Ethernet End-point, IP endpoint and Through modes.
- Support of hardware-assisted decoding of Precision Time Protocol (PTP) as defined in IEEE 1588-2008.
- Encapsulations: PTP over UDP over IPv4 as defined in IEEE 1588-2008 Annex D, PTP over IEEE 802.3 / Ethernet defined in IEEE 1588-2008 Annex F.

A.17.1. Results

- Presentation of peer clock details: Master identity, Grandmaster identity, Grand-master priority 1, Grandmaster priority 2, Grandmaster clock class, Grandmaster clock accuracy, Grandmaster clock variance, Grandmaster time source.
- TX and RX PTP frame counts classified by frame type.
- Sync Inter Arrival Delay (IAD) analysis: average and current.
- Sync Packet Total Delay (PTD): standard deviation, range.
- Sync Packet Delay Variation (PDV): current, maximum, average.
- Frequency offset between the master and the local clock (ppm).

A.18.Protocols

- *ARP* (IETF RFC 826).
- *DNS* (IETF RFC 1034, RFC 1035).
- *DHCP* (client side) (IETF RFC 2131).
- *Trace-route* application using UDP.
- *NTP* (client syde) (IETF RFC 5905) for system time synchronization with an external time server.

A.19.User Interface

- Direct configuration and management in graphical mode using the keyboard and display of the instrument.
- Remote access for configuration and management in graphical mode from remote IP site thought the Ethernet interface of the control panel.
- File management and download through web interface.

A.20.Platform

- Operation time with batteries (LiPO): 8 - 10 hours.
- Battery recharge time (LiPO): 4 hours.
- Operational range: -10ºC to +50ºC.
- Operation humidity: 10% - 90%.
- IP rating: 54.
- Configuration and report storage and export through attached USB port.
- TFT colour screen (480 x 272 pixels).
- Dimensions: 223 mm x 144 mm x 65 mm.
- Weight: 1.2 kg (with rubber boot).

Appendix B
Common Issues and Solutions

The following table summarizes some common configuration issues related with everyday use of Ether.Genius / Ether.Sync / Ether.Giga. Some of them are applicable only to a very specific situations but some others are valid as a general use advises. For example it is always better to use the local IP address as the source address for test traffic than an spoofed (manual) address whenever is possible.

Table B.1:

Issue	Solution
DHCP does not work	• Make sure that the network supports DHCP configuration and that your equipment is authorized to get a DHCP lease from the server. • In some situations Ether.Genius / Ether.Sync / Ether.Giga DHCP address configuration may be slow. The user can force immediate address negotiation by disabling DHCP for one moment and enabling it again.
Ping / Traceroute tests not available	• These are IP tests they are available only if the unit operation mode is set to IP endpoint mode.
eSAM, RFC 2544 or SLA test results are not available.	• eSAM, RFC 2544 or SLA analysis is active either in Port A results or in Port B results but never in both ports at the same time. Make sure that your test method settings (one-way or two way) are correct.
BER test results are not available	• BER testing require an special test payload. Configure the BERT payload in the *Payload* submenu within the port specific setup menu. Depending on your test setup you may need to configure the Port A and Port B payload to BERT.

Table B.1:

Issue	Solution
The unit is unable to establish link	• The test interface is configured in optical mode but you are using an electrical connector or the equipment is configured in electrical mode but you are using an electrical connector: Configure the right connector for your test by means the *Physical layer* settings in the port specific configuration menu. • You are using the optical interface but the light source is off. Please, remember that for security reasons the optical source is never enabled when the equipment is restarted. Make sure that the optical source is enabled in your optical tests. • The test interface does not support auto-negotiation but the equipment is configured to operate with auto-negotiation. The solution in this case is to disable auto-negotiation and force the bit rate to 10 or 100 Mb/s.
ARP does not resolve any destination MAC address.	• Your local proflie contains wrong settings. Your Ether.Genius / Ether.Sync / Ether.Giga unit decides where to send ARP requests using your local IP address, subnet mask and default gateway configuration. If there is an error with these settings your requests will be lost in the network or the destination will be unable to send the reply back to the origin. Make sure that your local profile contains the correct configuration. • ARP requests are encapsulated in a wrong frame structure. For example, the network may be unable to process standard Ethernet frames if it is waiting for VLAN tagged frames. Make sure that you are using the correct encapsulation. • ARP requests are delivered to unsupported VLANs. Make sure that you configure the correct VIDs in your frames.

Table B.1:

Issue	Solution
Test traffic is not transmitted through a switch	• Configuration for testing through a switch is more simple than for a router but you still need to check that you have configured the right destination MAC address, encapsulation and correct VLAN tags in the *Frame layer* menu. If you have configured a two-way test and in the remote end you have a traffic reflector, you have to configure the remote device MAC address as the destination address for your streams. If you use Port B as the traffic analyser in a one-way test, you have to configure the Port B MAC address as the destination address for your traffic streams.
Test traffic is not transmitted through a router	• You are using an incorrect MAC destination address. When you transmit data through a router the destination MAC address has to be configured to the router MAC address corresponding to the interface you are connected with. The easiest is to configure ARP in *Frame layer* settings and leave the tester auto-configure the right destination MAC address. • You are using a wrong frame encapsulation in your *Frame layer* settings. Make sure that you are using the correct encapsulation (VLAN, Q-in-Q, IEEE 802.1ad). • You are using incorrect IP settings in *Local profile* and *Network settings*. Your local profile should contain a unique IPv4 address. The default gateway must be placed in your local network (the network prefix should match the network prefix of your IP address). The remote address should be routable from the tester. For example you can not send traffic to a private address behind a NAT filter from the Internet. • The remote host should be able to reply to ARP requests. Note that Ether.Genius / Ether.Sync / Ether.Giga is unable to respond to ARP requests from addresses different to the local one. In case you need to receive test traffic in addresses different from the local one, you will need to configure static ARP entries in the router.

Table B.1:

Issue	Solution
No traffic detected in Filters panel.	• The traffic is being lost somewhere in the network. Use the global result statistics to know if this is the case. Make sure that the remote loopback device (if used) is properly configured. Check the *Network layer* (destination IP address), the *Frame layer* (destination MAC address, encapsulation, VLAN tags) and *Local profile* (IPv4 address, network mask, gateway and DNS server). The network will fail to send the test data to the destination if destination addresses and VID values are not configured to the right values. • You are receiving test traffic but filtering is not working properly. Use the *Match TX* setting in your filter configuration if you know that the network is not modifying the traffic in some way (CoS re-labelling, port modification in NAT filters or others). If the network is modifying the traffic, use a *Custom* filter instead.
All SLA test results are zero	• You may have a problem with the filter configuration. See how to solve problem with the filters (*No traffic detected in Filters*) • SLA results are computed only if you configure an special SLA payload in the generator and the analyser. for all your streams . You can set the SLA payload with the help of the *Payload* menu within the Port A setup. If you are using Port B as a traffic analyser, you will need to configure SLA payload in this port as well.
eSAM test fails to start	• If you are running a color-aware test the eSAM requires the green and yellow markers to have a different value for each service. For example, if you are using the VLAN priority bits as the color marker and you configure the green traffic priority bits to “1” in service number 1, then you have to configure the same field to something different to “1” for yellow traffic in the same service. Note that you can still configure the yellow traffic priority bits to “1” in services 2, 3 and 4.

Table B.1:

Issue	Solution
eSAM Policing test always fails.	• You are running a policing test but there is not a traffic admission mechanism configured in the network. The eSAM policing test is designed to test this mechanism and it fails if not conforming traffic is not dropped. If you want to avoid the eSAM test to fail for this reason you have to disable the policing test using the *Policing test* control in the *eSAM* menu.
eSAM or RFC 2544 test returns a *No traffic* message.	• If you are testing through a switched network, you may be using incorrect *Frame layer* settings. Check the solution on how to proceed when *Test traffic is not transmitted through a switch.* • If you are testing through a routed network, you may be using incorrect *Local profile* or *Network layer* settings. Check the solution about how to proceed when *Test traffic is not transmitted through a router.* • The network may be re-labelling traffic or modifying the test traffic in some way. Check the solution about how to proceed when *No traffic detected in Filters panel.*