



Frequent Asked Questions - Product technical

Chassis and Software

- [Which OS does the XenaManager™ run on?](#)
- [Is the Xena tester field upgradeable to future product releases?](#)
- [How do I change the chassis IP address?](#)
- [How do I connect to the Xena chassis if I have forgotten the chassis IP address?](#)
- [Why do the test port LEDs start blinking after power-on and reboot?](#)
- [Is the test module inside a XenaCompact identical to the XenaBay test modules?](#)
- [Are the XFP/SFP+/SFP transceiver plug-in modules hot-swappable?](#)
- [Why are the 10 Gbps XFP and SFP+ test modules factory equipped with optical transceiver plug-in modules?](#)
- [Can I daisy-chain multiple XenaCompact testers into a single logical test unit?](#)

Port Operation

- [How can I start/stop traffic on all ports simultaneously?](#)
- [Can multiple users reserve test ports on the same test module?](#)
- [Can multiple users use the same port on a test module?](#)
- [How can I move a port configuration between ports?](#)
- [Can I force a speed on a tri-speed 10/100/1000 port?](#)
- [Is there support for auto-negotiation on the test ports?](#)
- [Each time I restart traffic on a port, is the generated packet pattern 100% identical?](#)
- [Does a test port send any keep-alive packets?](#)
- [Can I view the port configuration of ports which are reserved by another user?](#)
- [Can I reserve a test port even when it is reserved by another user?](#)
- [What does it mean that the port property ppm speed-reduction is emulated?](#)

Traffic Generation

- [What is a stream?](#)
- [What is a test payload?](#)
- [How can I set the stream traffic rate to the maximum value?](#)
- [Why do streams have to be enabled?](#)
- [Is there support for stacked VLAN protocol \(Q-in-Q\)?](#)
- [How many streams can I configure?](#)
- [How does a modifier work?](#)
- [Can I specify a bursty traffic profile?](#)
- [Which packet sizes can be generated?](#)

- [Which types of packet payload can be inserted into the generated packets?](#)
- [Can I setup any packet header format?](#)
- [Is there support for Jumbo packets?](#)
- [Can the test ports answer ARP and PING requests?](#)
- [Can I use ARP to determine the Ethernet destination MAC address for IP packet streams?](#)
- [Where is the test payload inserted into the packet, and what is the format?](#)
- [What is a TID?](#)
- [Can the Xena tester generate UDP or TCP traffic?](#)
- [Is there support for invalid IFG spacing?](#)

Traffic Capture

- [What is traffic capture?](#)
- [Is Xena compatible with WireShark?](#)
- [How large is the packet capture buffer?](#)
- [Which packet protocols can the capture engine decode?](#)
- [How do I start/stop the traffic capture function?](#)
- [How can I define my packet capture criteria?](#)

Traffic Statistics

- [Which traffic statistics are provided?](#)
- [What is a filter?](#)
- [How is packet loss calculated?](#)
- [How can I analyze the inter-frame spacing between packets?](#)
- [How can I analyze the distribution of the packet latencies?](#)
- [Can I export traffic statistics to Excel?](#)
- [What is the accuracy of the latency measurements?](#)
- [Can I perform GPS timing synchronization between chassis for latency measurements?](#)
- [Can I setup an Rx filter for a 128-bit IPv6 address?](#)
- [How is packet payload integrity measured and checked?](#)
- [How do I calibrate latency measurements?](#)

Scripting and Automated Testing

- [Which features are available via scripting?](#)
- [Where can I find the documentation of the Xena API scripting commands?](#)
- [Is there a Tcl, Perl, or Java API available?](#)
- [How do I run a RFC2544 test](#)
- [Are there any wizards for setting up test patterns?](#)

Chassis and Software

Which OS does the XenaManager™ run on?

The XenaManager™ software runs on Windows XP (with .NET 2.0 or later) and Vista. It is a simple executable file, and does not require any installation.

[\[Back to top\]](#)

Is the Xena tester field upgradeable to future product releases?

Yes, the Xena testers are fully upgradeable to future product releases. Both the embedded software and the integrated test modules can be upgraded from the XenaManager™. If the upgrade includes a test module upgrade, the upgrade process requires a power cycle of the test unit.

[\[Back to top\]](#)

How do I change the chassis IP address?

The chassis management port is configured with a static IPv4 address (factory preset to 192.168.1.200). The management port's IP address can be changed from the XenaManager™. For detailed instructions, refer to our "Getting Started.pdf" document.

[\[Back to top\]](#)

How do I connect to the Xena chassis if I have forgotten the chassis IP address?

If you forget the IP address of the chassis you need another way to get in touch with the chassis. This is done by making a point-to-point connection from your PC to the "Ext" Ethernet port that is positioned to the right of the "Mgmt" port on the chassis front panel.

The "Ext" port is pre-configured with the following IP setup:

Address = 172.16.255.200

Subnet = 255.255.255.0

Gateway = none

You must configure your PC port statically to an IP address in the 172.16.255.x range, and then you will be able to PING the chassis again.

Now start the Manager, and connect to the chassis using address 172.16.255.200.

Now under CHASSIS PROPERTIES you can see which IP address is configured for the "Mgmt" port, and you can reserve the chassis and change it if necessary.

Note that the IP configuration of the "Ext" port cannot be changed, and that you should not configure the "Mgmt" port to use this subnet.

[\[Back to top\]](#)

Why do the test port LEDs start blinking after power-on and reboot?

Following reboot or power-on, the test port LEDs will start blinking after 2-3 minutes. This indicates that the tester is ready and now accepts XenaManager™ or scripting TCP/IP connections. When the first connection is made the LEDs revert to indicating link-sync status for each test port.

[\[Back to top\]](#)

Is the test module inside a XenaCompact identical to the XenaBay test modules?

Yes, a XenaCompact tester is functionally identical to a XenaBay tester.

[\[Back to top\]](#)

Are the XFP/SFP+/SFP transceiver plug-in modules hot-swappable?

The tester should be powered off when changing transceiver plug-in modules. Inserting or removing them with power on will not cause damage, but the tester will not detect any changes until it is re-booted or power-cycled.

[\[Back to top\]](#)

Why are the 10 Gbps XFP and SFP+ test modules factory equipped with optical transceiver plug-in modules?

All optical 10 Gbps test modules are sold equipped with optical transceiver modules. Xena's production and quality control is performed with the tester fully equipped with optical modules, to eliminate any possibility of errors caused by transceivers plug-in modules of poor quality.

Xena uses quality optical modules from the German manufacturer MergeOptics (formerly known as Infineon), and resells these optical modules at a very competitive price (at or below cost), when sold together with a Xena test module.

[\[Back to top\]](#)

Can I daisy-chain multiple XenaCompact testers into a single logical test unit?

No, the XenaCompact testers must be added one-by-one to the XenaManager™ testbed viewer, and script environments must setup TCP/IP connections to each individual tester.

[\[Back to top\]](#)

Port Operation

How can I start/stop traffic on all ports simultaneously?

Select the "Global" tab page in the XenaManager™ and click on the "START" button. This will start traffic on all ports which have a checkmark next to the "Start Traffic" button in their "TRANSMIT CONTROL" panel on the "Streams" tab page. Pressing "START" will also clear all TX/RX statistics on ports which have a checkmark next to the "Clear" buttons in their "Statistics" tab page. And it will also re-start capture and any histograms that are marked as under global control.

[\[Back to top\]](#)

Can multiple users reserve test ports on the same test module?

Yes, users can reserve test ports one-at-a time, and multiple users can therefore share the same test module.

[\[Back to top\]](#)

Can multiple users use the same port on a test module?

Any number of users can view the configuration of any port, as long as they have logged onto the chassis. In order to modify the configuration the port must be reserved.

[\[Back to top\]](#)

How can I move a port configuration between ports?

By using the port configuration save/load function. In the XenaManager™ testbed view, right click the port and select "Save port configuration". Then, right click the new port and select "Load port configuration". You can also save and load only the filter definitions for the port.

[\[Back to top\]](#)

Can I force a speed on a tri-speed 10/100/1000 port?

Yes, the tri-speed 10/100/1000M electrical test ports can be forced into a fixed speed mode of either 10 Mbps, 100 Mbps, or 1000 Mbps, as opposed to operating in standard auto-negotiation speed mode.

[\[Back to top\]](#)

Is there support for auto-negotiation on the test ports?

Auto-negotiation is only provided on the tri-speed 10/100/1000M electrical test ports.

[\[Back to top\]](#)

Each time I restart traffic on a port, is the generated packet pattern 100% identical?

Yes, the generated traffic pattern is 100% identical each time the traffic is re-started, including all pseudo random generated numbers. To make each run produce a different pattern, you should clear the "Random seed" field in the "PORT PROPERTIES" panel on the "Properties" tab page.

[\[Back to top\]](#)

Does a test port send any keep-alive packets?

The test ports do not generate any keep-alive packets, and will only generate traffic which is configured and specified by the user. A test port which is not sending any traffic into the DUT/SUT, but is receiving traffic from the DUT/SUT, should therefore be configured to generate keep-alive packets e.g. at a rate of 1 packet per second to avoid MAC table timeout in the DUT/SUT. You can also send a single packet by pressing the "Transmit" button in the "PORT PROPERTIES" panel.

[\[Back to top\]](#)

Can I view the port configuration of ports which are reserved by another user?

Yes, any user can view all configuration and traffic statistics for any test port, also when the test port is reserved by another user. To view the latest port stream configuration of a port reserved by another user, the user must manually select the port, right-click the mouse, and select "Refresh port configuration".

[\[Back to top\]](#)

Can I reserve a test port even when it is reserved by another user?

Yes, you can reserve a port even when it is reserved by another user, by performing a port "Relinquish" command first on the "PORT PROPERTIES" panel. The relinquish command will release the current port reservation, and the port can now be reserved by you.

[\[Back to top\]](#)

What does it mean that the port property ppm speed-reduction is emulated?

When the test port ppm speed-reduction function is enabled by specifying a ppm speed-reduction greater than zero, the system will automatically insert an idle traffic period every second, where the length of the idle period corresponds to the specified speed-reduction specified in units of ppm. The resulting effective test port speed is shown below the speed-reduction value. Enabling the ppm speed-reduction function does not change the clock speed of the test port link.

[\[Back to top\]](#)

Traffic Generation

What is a stream?

Streams are the basic mechanism for generating outgoing traffic on a port. The configuration of a stream is divided into two areas:

- *Transmission profile*, which specifies how the packets should be spaced in time. The primary parameter is the overall *rate*, which can be specified as an absolute bit rate, as a percentage of the port's overall rate, or as the number of packets per second. In addition, the train of packets can be made more or less bursty.
- *Packet content*, which specifies the actual data bytes comprising each packet. The content is auto-generated by the modules, using a common packet header and a varying-size payload. The packet

header can be specified according to a number of common protocols. You can apply modifiers to individual fields of the packet, which will then be varied on a packet-by-packet basis, allowing a multitude of separate flows to be generated by a single stream.

[\[Back to top\]](#)

What is a test payload?

Each packet of a stream can be tagged with a special test payload, which contains time-stamp and sequence number information. This allows the receiving port, which may be on a separate chassis and indeed in a separate physical location, to verify proper transmission of every packet of the stream. Part of the test payload is a test payload id, or TID, which allows the receiving port to differentiate between multiple interleaved packet streams, possibly originating from different ports and different chassis. Test payloads form the basis for the traffic integrity checking and latency measurements.

[\[Back to top\]](#)

How can I set the stream traffic rate to the maximum value?

If you configure a stream rate which exceeded the maximum allowed rate, the rate is automatically capped to the maximum allowed value. The maximum allowed value takes into account the test port speed, and the configured rate of any other stream enabled on the same test port.

[\[Back to top\]](#)

Why do streams have to be enabled?

You can define more streams for a port than can be transmitted at one time. This allows you to have a variety of streams defined, and enabling them in different combinations. You can even set the rate of a (non-enabled) stream to more than 100% which allows you to play around with the percentage/bps/pps values before enabling the stream.

[\[Back to top\]](#)

Is there support for stacked VLAN protocol (Q-in-Q)?

Yes, Q-in-Q packets can be generated by simply inserting two VLAN segments in the stream packet header. Packet headers with three or more stacked VLANs can also be configured.

[\[Back to top\]](#)

How many streams can I configure?

On the transmit side, up to 256 streams can be configured per 10 Gbps test port, and up to 32 streams can be configured per 10/100/1000M test port. The insertion of the test payload (TID) is optional per stream. You can use field modifiers to make multiple flows within a single stream. On the receive side, a 10 Gbps test port can measure throughput and performance statistics on up to 2048 simultaneous incoming streams, and a 10/100/1000M test port can measure statistics on up to 680 simultaneous incoming streams.

[\[Back to top\]](#)

How does a modifier work?

You can apply modifiers to individual fields of the packet headers, which are then varied on a packet-by-packet basis as the packets for a particular stream are generated and transmitted. Modifiers are typically used on MAC addresses, VLAN ids, IP addresses, the IP QoS field, and UDP ports.

Two modifiers can be configured per stream. The position of the modifier can be located anywhere within the first 256 bytes of the packet, and the size of the modifier can be from 1 to 16 bits. The modifier can traverse a specified min-to-max range by either incrementing (in steps of 1-16), decrementing (in steps of 1-16), or as a random value. The two modifiers can operate independently of each other, or one can be made to change slower than the other, thereby allowing for traversing two-dimensional ranges of packet header fields.

[\[Back to top\]](#)

Can I specify a bursty traffic profile?

Yes, each stream can define its own unique traffic burstiness profile. The traffic profile is defined by the number of packets per burst, and the burst density which defines the spacing between the packets within the burst. The density can range from fully uniform (even spacing between all packets) to fully bursty (all packets within a burst are transmitted back-to-back).

[\[Back to top\]](#)

Which packet sizes can be generated?

Standard Ethernet packet sizes, invalid packet sizes, and jumbo size packets can be generated, with the supported packet size range being the full range of 56-16383 bytes. Per stream, the packet sizes can be fixed, incrementing, decrementing, random, or 'butterfly' (min, max, min+1, max-1, min+2, etc) over a defined range.

[\[Back to top\]](#)

Which types of packet payload can be inserted into the generated packets?

Two types of packet payload can be defined for a stream: either a sequence of incrementing bytes up through the packet, or a custom defined 1-18 byte pattern which is repeated up through the packet.

[\[Back to top\]](#)

Can I setup any packet header format?

Yes. The XenaManager™ currently supports the following protocols: Ethernet, Ethernet II, VLAN, ARP, IPv4, IPv6, UDP, TCP, LLC, SNAP, GTP, ICMP, RTP, RTCP, and STP. You can also specify packet header content byte-by-byte using the hex editor.

[\[Back to top\]](#)

Is there support for Jumbo packets?

Yes, packet sizes up to 16383 bytes can be generated per stream.

[\[Back to top\]](#)

Can the test ports answer ARP and PING requests?

Yes. On the "PORT PROPERTIES" panel you can specify an IP address for the port, and whether the port should respond to incoming ARP and PING requests for this address.

[\[Back to top\]](#)

Can I use ARP to determine the Ethernet destination MAC address for IP packet streams?

Yes. In a stream definition containing an IP packet header, you can use ARP to send a request on the test port for the corresponding MAC address, which is then inserted into the packet header.

[\[Back to top\]](#)

Where is the test payload inserted into the packet, and what is the format?

The test payload is a 20 byte data field, which is inserted at the end of the packet just before the Ethernet FCS field. The test payload contains a timestamp, a packet sequence number, a TID, a payload integrity check value, and a long checksum which the receiving test port uses to identify if the incoming packets carry a Xena test payload.

[\[Back to top\]](#)

What is a TID?

A TID is a field in the test payload used to identify streams on the receiving side. It is a value in the range from 0 to 0xFFFF. TID values must be unique for all streams terminating at a particular port. The easiest way to accomplish this is simply to use unique TID values for every stream definition across the whole testbed.

When the XenaManager™ can identify a unique source for a particular received TID it can calculate packet loss by subtracting the TX and RX packet counts.

[\[Back to top\]](#)

Can the Xena tester generate UDP or TCP traffic?

Yes, traffic streams can be configured with UDP headers. For TCP, only stateless TCP packets can be generated, i.e. there is no support for stateful TCP traffic.

[\[Back to top\]](#)

Is there support for invalid IFG spacing?

Yes, the packet spacing can be specified down to 16 bytes (where 20 bytes is the standard), where the Inter-Frame Gap is reduced to 8 bytes (where 12 bytes is the standard). The Ethernet Preamble is always 8 bytes.

[\[Back to top\]](#)

Traffic Capture

What is traffic capture?

Capture allows selected packets of the incoming traffic on a port to be retained and inspected.

- *Triggers* specify when to start capturing, when to stop, and which packets to keep while capturing is active. These events can be specified as user-defined filters or as built-in events such as packet checksum errors.
- Results can be viewed directly within the Xena environment, and packets can be saved for further analysis by external tools such as WireShark.
- There is also a packet-by-packet graphical display of key capture parameters, such as packet length and inter-frame gap.

[\[Back to top\]](#)

Is Xena compatible with WireShark?

Yes, the captured packets can be saved in WireShark “.pcap” format, for further analysis in the WireShark environment. This is useful when analyzing packet protocols which are not (yet) supported in the XenaManager™.

[\[Back to top\]](#)

How large is the packet Capture buffer?

The wire-speed capture buffer is 64KB for the 10 Gbps test port, and 16KB per 10/100/1000M test port. The number of captured bytes per packets can be limited, e.g. to 64 bytes, so that only the packet header is captured. Then approximately 1000 or 250 packets can be captured per 10 Gbps and 10/100/1000M test port respectively. If the captured criteria is specified to match an incoming lower speed stream with a 2-5 Mbps rate (such as a VoIP or IPTV stream), the capture buffer can store 4096 packets at any packet length.

[\[Back to top\]](#)

Which packet protocols can the capture engine decode?

Ethernet, Ethernet II, VLAN, ARP, IPv4, IPv6, UDP, TCP, LLC, SNAP, GTP, ICMP, RTP, RTCP, and STP. For analyzing other protocols, the captured data must be exported to WireShark.

[\[Back to top\]](#)

How do I start/stop the traffic capture function?

The capture start/stop trigger events can be the arrival of a packet which matches a specific user-defined packet header filter (such as a specific Ethernet/IP address, VLAN ID, IP TOS etc) or a packet FCS checksum error.

[\[Back to top\]](#)

How can I define my packet capture criteria?

The capture function can be configured to retain packets which match a user-defined packet header filter (such as a specific Ethernet/IP address, VLAN ID, IP TOS etc), all traffic, or all traffic which does not include a test payload.

[\[Back to top\]](#)

Traffic Statistics

Which traffic statistics are provided?

Statistics provide basic measurements of the amount and rate of various packet flows, at the bit/byte level and at the packet level:

- Per port, the overall amount of traffic sent and received on the port, with and without test payloads.
- Per stream, the amount of traffic transmitted for each outgoing stream defined for the port.
- Per test payload, the amount of traffic received containing test payloads with each particular id value. Also provides sequence, disorder, payload integrity error counts and latency measurements.
- Per filter, the amount of traffic conforming to each of the user-defined filters.

[\[Back to top\]](#)

What is a filter?

Filters act on the incoming traffic on a port, and recognize packets based on user-defined patterns. Filtering is based on a two-level approach:

- *Terms*, which specify simple true/false properties of a packet. *Match terms* look for a particular bit pattern at a particular position, and *length terms* look for packets that are longer or shorter than a given threshold.
- *Conditions*, which are compound Boolean expressions involving one or more terms. The expression can use the logical operators &, | and ~ in any combination.

Traffic statistics counters are automatically provided for every filter defined by the user, in the "RECEIVE STATISTICS" panel on the "Statistics" tab page. The filter can also be used to start/stop packet capture, and select which packets must be retained.

[\[Back to top\]](#)

How is packet loss calculated?

The XenaManager™ automatically calculate loss for traffic streams which include a test payload and where a unique source for the TID can be identified. The packet loss is calculated when the stream is no longer transmitting packets, and it is the difference between the number of transmitted packet and the number of received packets, for the particular stream. In contrast, the sequence and integrity errors are reported on-the-fly as they are decoded exclusively on the receive side.

[\[Back to top\]](#)

How can I analyze the inter-frame spacing between packets?

The “Histograms” tab page allows you to configure to display the IFG in histogram format, real-time, as packets arrive, classified and counted according to the IFG value. Alternatively, the packet capture function displays the packet-by-packet IFG spacing in graphical format. Histograms can also classify and count packets based on their length.

[\[Back to top\]](#)

How can I analyze the distribution of the packet latencies?

The “Histograms” tab page allows you to display the latency in histogram format, real-time as packets arrive, where the histogram bars represent latency bins, e.g. 0-10 ms, 10-20ms, 20-30ms, etc.

[\[Back to top\]](#)

Can I export traffic statistics to Excel?

Yes, traffic statistics can be exported in “.csv” format, which can be read into Excel.

[\[Back to top\]](#)

What is the accuracy of the latency measurements?

The latency measurement accuracy is 16 ns for the 10 Gbps test port, and 32 ns for the 10/100/1000M test ports. Latency can only be measured between TX and RX test ports located in the same chassis.

[\[Back to top\]](#)

Can I perform GPS timing synchronization between chassis for latency measurements?

No, GPS based timing synchronization between chassis is a roadmap feature. Latency can therefore only be measured between TX and RX test ports located in the same chassis.

[\[Back to top\]](#)

Can I setup an RX filter for a 128-bit IPv6 address?

Yes. A total of six 64-bit match terms are available per test port, so an IPv6 address filter will therefore require the allocation of two match terms for this purpose.

[\[Back to top\]](#)

How is packet payload integrity measured and checked?

As packets traverse networks, IP header contents may change - resulting in recalculation of packet CRC values. To validate network device correctness, the data integrity function allows packet payload contents to be verified, to test and verify that the payload is not disturbed as the network device changes header fields.

When measuring packet payload integrity, the packet payload pattern must be set to an incrementing pattern. A check-value in the test payload identifies the packet byte offset and value from which the incrementing payload must be present. If the RX test port does not recognize an error free incrementing test payload pattern from this offset position and up until the last packet payload byte, a payload integrity error is reported.

[\[Back to top\]](#)

How do I calibrate latency measurements?

Connect the TX and RX test ports between which latency is measured, with a short loop-cable, transmit one or more packets with test payloads between the two ports, and then perform the “Calibrate” function on the “Statistics” tab page. Then, remove the loop-cable and run the traffic test. The measured latencies will now only include the latency of the DUT/SUT.

[\[Back to top\]](#)

Scripting and Automated Testing

Which features are available via scripting?

The full feature set is available via the TCP/IP based command line interface. Anything that can be done from the XenaManager™ can also be done via text-based commands.

[\[Back to top\]](#)

Where can I find the documentation of the Xena API scripting commands?

Refer to the *Xena Scripting Specification* - the programmer's reference manual for the scripting capability. The manual is available to registered Xena customers, or under NDA.

[\[Back to top\]](#)

Is there a Tcl, Perl or Java API available?

No, there are not yet any generic scripting API wrappers available, but Xena customers have implemented their own proprietary Tcl, Perl, and Java scripting wrappers, where the wrapper converts from Xena's generic TCP/IP command line interface, to the customer's own and proprietary test equipment Tcl/Perl/Java API.

[\[Back to top\]](#)

How do I run a RFC2544 test?

By using Xena's RFC2544 implementation, provided in the form of an Excel spreadsheet programmed in Visual Basic for Applications (VBA).

[\[Back to top\]](#)

Are there any wizards for setting up test patterns, e.g. a full mesh traffic load on an n-port DUT?

Not yet, a number of wizards are however defined as roadmap for the Xena platform, and these Wizards will be provided in the form of Excel spreadsheets programmed with Visual Basic for Applications (VBA).

[\[Back to top\]](#)