

Scheduling modes



APPLICATION NOTE

Merging multiple Xena Traffic Generator streams onto one physical port.

A Traffic Generator from Xena Networks provides very flexible options for generation of Ethernet test traffic. The traffic is structured as a number of individually configurable traffic streams. The Xena Traffic Generator supports four ways (scheduling modes) of merging the streams into the test traffic that is sent from one physical port. This application note explains the four scheduling modes.

Contents

| | |
|--|----|
| Application Note..... | 3 |
| Scheduling Modes | 3 |
| Normal Interleaved Scheduling..... | 4 |
| Strict Uniform Scheduling | 5 |
| Sequential Scheduling | 5 |
| Burst Scheduling..... | 6 |
| Understanding How Gaps Are Aligned | 8 |
| Inspecting Burst Settings..... | 9 |
| Preview Traffic..... | 9 |
| Timing Source | 11 |
| Stagger Traffic Generation on Ports..... | 11 |

The Xena Networks Layer 2-3 testers – the XenaBay and XenaCompact – allow advanced, flexible generation of Ethernet test traffic from 100 Gbps to 100 Mbps depending on installed test modules. The traffic is structured as a number of individually configurable traffic streams providing very flexible options for generation of Ethernet test traffic. The way these streams are merged together can also be configured to match different test applications. The Xena Traffic Generator supports four ways (scheduling modes) of merging the streams into the test traffic that is sent from one physical port.

This application note explains how the four scheduling modes works, and how you can use these scheduling modes to generate test traffic matching your application.



Figure 1: The versatile and powerful Xena Networks Layer 2-3 testers XenaBay and XenaCompact

The XenaBay and XenaCompact can be controlled through the XenaManager, an easy-to-use application running on a Windows PC. XenaBay and XenaCompact can also be controlled by scripting through a Command Line Interface (CLI) for automation applications. Everything you can do with XenaManager can also be done via scripting, using simple CLI text commands. The examples in this application note will be illustrated with XenaManager screen shots.

SCHEDULING MODES

The scheduling mode for outgoing traffic from Xena Networks Traffic Generator port, specifies how multiple logical streams are merged onto one physical port. There are four scheduling modes:

- Normal interleaved
- Strict Uniform
- Sequential
- Burst (requires SW release 76 or higher)

The Scheduling Mode is selected as the Port TX Mode in the TX Profile section of the port properties page in XenaManager as shown in figure 2.

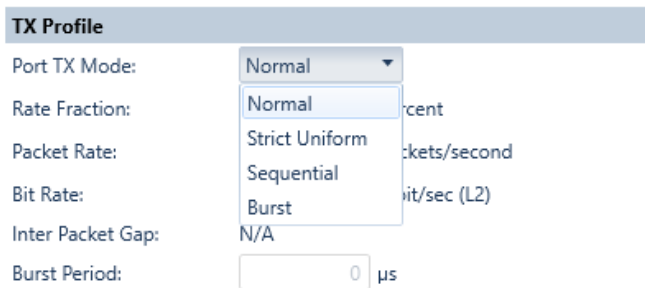


Figure 2: The TX Profile section of the port properties page

NORMAL INTERLEAVED SCHEDULING

In the Normal Interleaved scheduling mode (which is the default scheduling mode) streams are treated independently and are merged into a combined traffic pattern for the port, which honors each stream's ideal packet placements as well as possible. If all streams have identical settings – including the same fixed packet size – the packets will be distributed equally as illustrated in figure 3.



Figure 3: Distribution of test packets from 4 streams with identical settings



Figure 4: Distribution of test packets from 4 streams with different frame size settings

If the frame size is different from between the streams, the distribution of the packet in the transmitted test traffic will change as illustrated in figure 4: In this case the fixed frame size in stream 0 is 64 bytes, while it is 128 bytes in streams 1, 2 and 3; all streams have same rate fraction. Hereby stream 0 packets will appear more often in the merged stream than packets from streams 1, 2 and 3.

For some applications it will be relevant that the test packets from one or more streams are grouped together in bursts. This is controlled in the Transmission Profile section of the Stream Properties page in XenaManager as shown in figure 5. The Burst size parameter identifies the number of packets in the burst, while the burst density identifies how close the packets are together in the burst – 100 % density is the closest possible. The burst signature in figure 5 indicates the gap between the packets within the burst. Figure 5 shows how two different burst densities (100 % and 20 %) are illustrated. The “Inter Packet Gap” line (in green boxes) shows the calculated value of the distance (in ns and bytes) between the packets in the burst. The “Inter Burst Gap” line (in blue boxes) shows the calculated value of the distance (in ns and bytes) between last packet in one burst and the first packet in the next burst. The packets in a burst may still be merged with packets from other active streams during the burst, but the majority of the packet will come from the burst. Outside the burst only packets from other active streams will be sent.




Figure 5: Stream Transmission profile with parameters for setting of bursts in the Normal Scheduling mode

STRICT UNIFORM SCHEDULING

The Strict Uniform scheduling mode is a slight variation of Normal Interleaved scheduling mode. The Strict Uniform scheduling mode emphasizes strict uniformity of the inter-packet-gaps as more important than hitting the stream rates completely accurate.

SEQUENTIAL SCHEDULING

In the Sequential scheduling mode each stream in turn contribute one or more packets, before continuing to the next stream, in a cyclical pattern. The Sequential scheduling mode is useful if a specific sequence of test packets with different characteristics (e.g. regarding size or contents) is needed for the test traffic.

The number of packets that a stream contributes is entered in the Seq.Packets field in the Transmission Profile for the stream (see figure 6). The individual stream rates are ignored; the overall rate is determined at the port-level in the Port TX Profile section (figure 6). This in turn determines the rates for each stream, considering their packet lengths and counts. The maximum number of packets in a cycle (i.e. the sum of Seq.Packets for all enabled streams) is 500.

| TX Profile | | Transmission Profile | |
|-------------------|------------------------|----------------------|------------------------|
| Port TX Mode: | Sequential | Rate Fraction: | 10.000 percent |
| Rate Fraction: | 10.00 percent | Packet Rate: | 1488095 packets/second |
| Packet Rate: | 1488095 packets/second | Bit Rate L2: | 761.9048 Mbit/sec |
| Bit Rate: | 761.905 Mbit/sec (L2) | Bit Rate L1: | 1000.0000 Mbit/sec |
| Inter Packet Gap: | 620 ns (776 bytes) | Rate Cap: | Cap Rate |
| Burst Period: | 0 μs | Inter Packet Gap: | N/A |
| | | Seq.Packets: | 1 packets |
| | | Burst Size: | 10 packets |
| | | Burst Density: | 20 percent |
| | | Inter Packet Gap: | 0 bytes |
| | | Inter Burst Gap: | 0 bytes |
| | | Inter Burst Gap: | N/A |

In Sequential Scheduling mode the transmitted rate is determined by the Rate Fraction, Packet Rate and Bit Rate parameters, defined at port level

Port TX Profile

Stream Transmission Profile

Figure 6: Profile definitions in the Sequential scheduling mode




Figure 7: Sequential test packet flow examples with different Seq.Packets settings

Figure 7 illustrates how a traffic pattern is affected by different settings of Seq.Packets. At a first sight, example 1 in figure 7 looks the same as the example in figure 3 (where Normal Interleaved scheduling is used). Please observe however that example 1 in figure 7 will look the same regardless of the frame sizes defined in the different streams. As shown in figure 4 the distribution of frames from the involved streams will depend on the frame sizes defined in the streams when Normal Interleaved scheduling is used.

BURST SCHEDULING

SW release 76 (or higher) include the Burst scheduling mode. In this mode, frames from the streams on a port are sent as bursts as depicted in figure 8. During a burst, only packets from one stream will be sent.




Figure 8: Bursts from 2 streams in Burst scheduling mode

The Burst Period is defined by the Burst Period parameter (see figure 9) in the TX Profile section of the Port Properties page. The individual bursts are defined in the Transmission Profile section of the Stream Properties page (see figure 9):

- The number of packets in a burst is defined by the stream Burst Size parameter
- The Inter Packet Gap (in bytes) is defined by the stream Inter Packet Gap parameter together with the Minimum Inter-Frame Gap (see figure 11): The two values are added
- The Inter Burst Gap (in bytes) is defined by the stream Inter Burst Gap parameter together with the Minimum Inter-Frame Gap (see figure 11): The two values are added

| TX Profile | | Transmission Profile | |
|----------------------|------------------|--------------------------|------------------------|
| Port TX Mode: | Burst | Rate Fraction: | 10.000 percent |
| Rate Fraction: | 0 percent | Packet Rate: | 3720238 packets/second |
| Packet Rate: | 0 packets/second | Bit Rate L2: | 1904.7619 Mbit/sec |
| Bit Rate: | 0 Mbit/sec (L2) | Bit Rate L1: | 2500.0000 Mbit/sec |
| Inter Packet Gap: | N/A | Rate Cap: | Cap Rate |
| Burst Period: | 100 μs | Inter Packet Gap: | N/A |
| | | Stop After: | 0 packets |
| | | Burst Size: | 100 packets |
| | | Burst Density: | 100 percent |
| | | Inter Packet Gap: | 10 bytes |
| | | Inter Burst Gap: | 100 bytes |
| | | Inter Burst Gap: | N/A |

Port TX Profile

Transmission Profile in Stream Properties page

Figure 9: Profile definitions in the Burst scheduling mode

A test module will support less streams in Burst scheduling mode than in Normal mode. Number of streams supported depend on the Burst Size and Gap sizes defined for the streams. An error message will appear when traffic generation is started if too many streams are activated on a port in Burst mode.

UNDERSTANDING HOW GAPS ARE ALIGNED

Ethernet frames sent on a line are separated by Inter Frame Gaps (IFG). Prior to each frame a 7 byte Preamble and a Start of Frame Delimiter (SFD) byte will indicate the start of a new frame as depicted in figure 10 for an Ethernet/IP frame with the default value for the IFG.




Figure 10: Ethernet/IP frame format

For simplicity the IFG, Preamble and SFD are presented as one parameter (Inter Packet Gap or Inter-Frame Gap) in the XenaManager. The minimum value will default to 20 bytes but can be set to any value between 16 and 63 bytes (some modules support values between 18 and 63 bytes) in the Layer-1 Control part of the Port Properties page as shown in figure 11.




Figure 11: Minimum Inter-Frame Gap is set to 20

The Inter-Frame Gap defined in figure 11 will be the minimum Inter-Frame Gap for all packet generated in all streams on the port. This means that this value is added to the Inter Packet Gap and Inter Burst Gap parameters (figure 9). Therefore, if the Inter Packet Gap is set to 10 as shown in figure 9 and the Minimum Inter-Frame Gap is set to 20 (figure 11), the total value that will be used as Gap will be $10+20 = 30$ bytes.

In the Ethernet standard it is defined how the Ethernet frame is built. The standard is based on packets being aligned to boundaries of a fixed number of bytes. The Inter Packet Gap is adjusted to achieve the alignment with the boundary:

- At 40G and higher speeds the boundary size is 8 bytes, meaning that the actual Inter Packet Gap will toggle between values that are multiples of 8, just below and above the target value. In the example above, where 30 bytes is the target value, the actual Inter Packet Gap at 100G will toggle between 24 and 32; over time the average size of all Inter Packet Gaps will be 30 bytes.
- At 10G and 25G the boundary size is 4, meaning that actual Inter Packet Gap will toggle between values that are multiples of 4. In the example above, where 30 bytes is the target value, the actual Inter Packet Gap at 10G will toggle between 28 and 32; over time the average size of all Inter Packet Gaps will be 30 bytes.
- At 1G the boundary size is 2, meaning that actual Inter Packet Gap will toggle between values that are multiples of 2.

INSPECTING BURST SETTINGS

When a burst sequence has been configured in the Burst scheduling mode, it may be useful to inspect that the packet bursts being transmitted are as expected. A quick way to achieve this is to use the preview facility of XenaManager. This will activate a Tx-to-Rx loopback for the port, start the traffic generation and activate the XenaManager capture function. Table 1 and figure 12 explains the Xena Traffic Generator loopback modes.

| Loopback mode | Explanation |
|----------------------|---|
| Off | Traffic flows naturally out of the port |
| L1 RX-to-TX | Any received packet is bounced back through TX |
| L2 RX-to-TX | Same as above yet it also swaps MAC SRC<->DST (Not available on some test modules) |
| L3 RX-to-TX | Same as above yet it also swaps IP SRC<->DST (Not available on some test modules) |
| TX(on)-to-RX | Packet goes out of TX but also internally direct to RX (Not available at 2.5G and 5G rates) |
| TX(off)-to-TX | Packet goes directly to RX (No link sync needed) (Not available at 2.5G and 5G rates) |
| Port-to-Port | Any received packet goes out through the neighbor port (Not available on some test modules) |

Table 1: Xena Traffic Generator loopback mode

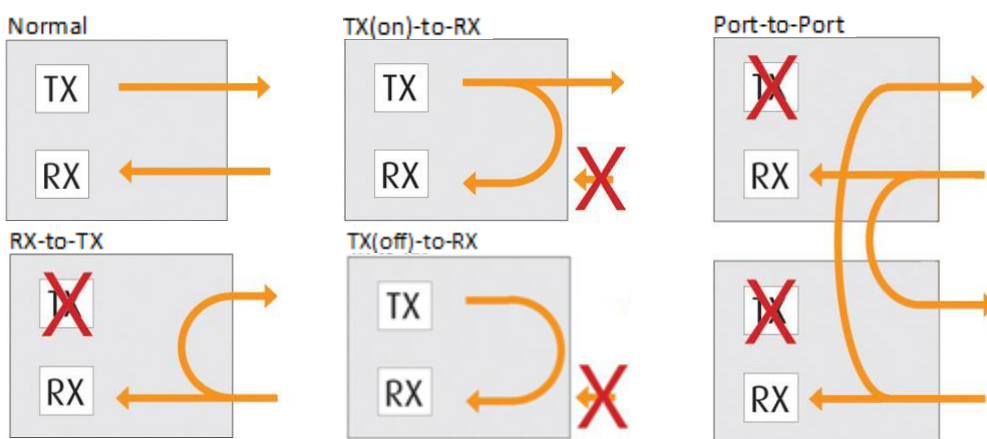


Figure 12: Xena Traffic Generator loopback modes

PREVIEW TRAFFIC

You can activate the preview facility of XenaManager by right-clicking on the parent port for the burst sequence and selecting **Preview Port Traffic** from the menu that appears. A pop-up will explain the process after Start Preview is pressed (figure 13).

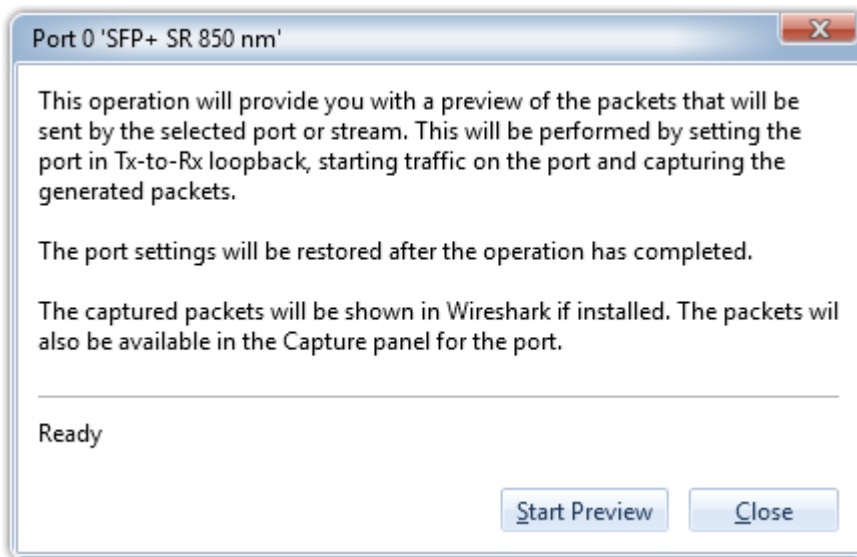


Figure 13: Preview pop-up

When Start Preview is pressed, the traffic will run until the capture buffer runs full or for 10 seconds if the buffer is not filled up in that period. When the packets have been captured, they can be inspected in the XenaManager Capture panel and in Wireshark if that is installed on the PC running the XenaManager. Figures 14 and 15 show the Capture panel results for bursts at 100G and 10G where the IPG is set to 10 bytes, which together with the basic 20 byte Inter-Frame Gap give a 30 byte gap that should be expected in average between the captured frames. The IFG column in the Capture panel shows the gap in bytes between a given frame and the previous frame. IFG for the first captured frame in figures 14 and 15 is 0, as there is no previous frame captured. Over 4 consecutive frames @ 100G and 2 consecutive frames @ 10G the average value is 30, i.e. the expected value.

| Capture Result Grid | | Capture Graph | | | | | | |
|---------------------|----------------|-------------------------|-------------|--|-------------------|------------------------|-------------|------------|
| Capture Results | | Captured Packets: 2,745 | | <input checked="" type="checkbox"/> Time Relative To Start | | | | |
| # | Timestamp (ns) | Latency (ns) | IFG (bytes) | Source | Destination | Protocol | Full Length | Cap Length |
| 0 | 1,516,048 | 0 | 0 | 04:F4:BC:44:78:30 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 1 | 1,516,056 | 0 | 24 | 04:F4:BC:44:78:30 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 2 | 1,516,064 | 0 | 32 | 04:F4:BC:44:78:30 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 3 | 1,516,072 | 0 | 32 | 04:F4:BC:44:78:30 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 4 | 1,516,072 | 0 | 32 | 04:F4:BC:44:78:30 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 5 | 1,516,080 | 0 | 24 | 04:F4:BC:44:78:30 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |

Figure 14: Capture result at 100G

| Capture Result Grid | | Capture Graph | | | | | | |
|------------------------|----------------|--|-------------|-------------------|-------------------|------------------------|-------------|------------|
| Capture Results | | | | | | | | |
| Captured Packets: 989 | | <input checked="" type="checkbox"/> Time Relative To Start | | | | | | |
| # | Timestamp (ns) | Latency (ns) | IFG (bytes) | Source | Destination | Protocol | Full Length | Cap Length |
| 0 | 318,480 | 0 | 0 | 04:F4:BC:28:0B:80 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 1 | 318,544 | 0 | 28 | 04:F4:BC:28:0B:80 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 2 | 318,624 | 0 | 32 | 04:F4:BC:28:0B:80 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 3 | 318,704 | 0 | 28 | 04:F4:BC:28:0B:80 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 4 | 318,776 | 0 | 32 | 04:F4:BC:28:0B:80 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |
| 5 | 318,848 | 0 | 28 | 04:F4:BC:28:0B:80 | 00:00:00:00:00:00 | ETHERNET/IP/Raw/XENA_T | 64 | 64 |

Figure 15: Capture result at 10G

TIMING SOURCE

As you can see in figures 14 and 15 a time stamp is assigned to every captured frame. The best accuracy for the timing of the packets in the bursts will be achieved if the Timing Source is set to “Local Module Time” in the Timing Configuration section of the Module Properties page (see figure 16).



Figure 16: Timing Source selection

STAGGER TRAFFIC GENERATION ON PORTS

For some applications it can be relevant to stagger/delay start of traffic generation on one port relative to other ports. This can be achieved with the “Stagger Factor” facility of the Xena Traffic generator, which delays start of traffic generation on one port relative to activation of global start. The delay is programmed in steps of 64 μsec in the Layer 1 Control section of the Port Properties page and can have values between 0 and 31250. If the maximum value is used (see figure 17) the traffic generation of the port is delayed $31250 * 64 \mu\text{sec} = 2 \text{ seconds}$.

| Layer-1 Control | |
|--------------------------|---|
| Port Speed Selection: | <input type="text"/> |
| Min. Inter-Frame Gap: | <input type="text" value="20"/> bytes |
| Speed Reduction: | <input type="text" value="0"/> ppm, emulated |
| Current Port Speed: | 10 Gbit/s |
| Effective Port Speed: | 10 Gbit/s |
| Auto-Negotiation Enable: | <input type="checkbox"/> |
| Stagger Factor: | <input type="text" value="31250"/> x 64 μ s |

Figure 17: Setting of Stagger Factor

The Stagger Factor will work between ports on test modules installed in the same chassis.