



Testing at 100% load with line rate speed compensation

December, 2009

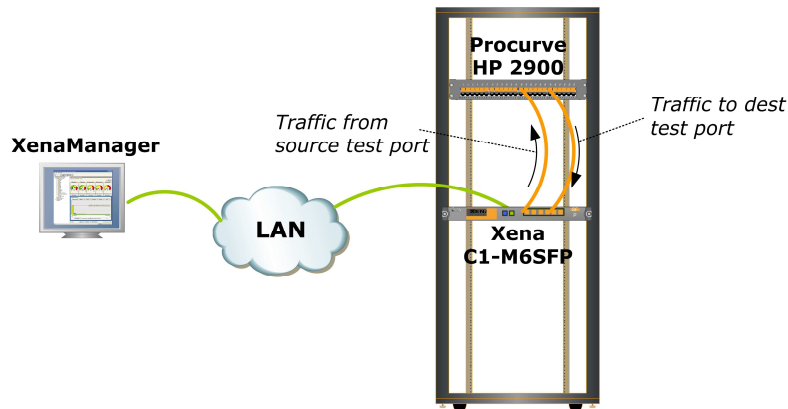
Application Note

The Xena testers can verify traffic forwarding performance, protocol scalability and services delivering capabilities of switching and routing devices across the enterprise, metro/edge and core.

Prior to stress testing the forwarding capability, special consideration must be applied to eliminate potential packet loss caused by small differences in line rates between the Xena test ports and the network ports of the Device-Under-Test (DUT). This application note uses an example test setup with a HP Procurve 2900 24G switch to illustrate how to compensate for small differences in line rates between the Xena tester and the DUT.

Test Setup

The test setup is shown in Figure 1. A source test port sends Ethernet frames through the device under test (DUT) while real-time performance statistics are measured by the Xena tester on the destination port. In this test the DUT is a HP Procurve 2900 24G (J9049A) Switch.



▪ Figure 1 Test setup with HP Procurve 2900 24G

Compensate for Line Rate Clock Differences

The first step prior to performing any stress load testing is to ensure that small differences in line rates will not introduce unwanted packet loss due to buffer overflow inside the DUT and lack of a congestion control backpressure mechanism from the DUT to the Xena tester, since that can easily cause confusion when analyzing the forwarding capabilities of the DUT.

The Ethernet line rate clock frequency is typically defined with an accuracy of +/- 100 parts per million (ppm). In a worst case scenario, the line rate frequency of the Xena tester is 100 ppm *above* the nominal frequency while the line rate frequency of the DUT is 100 ppm *below* the nominal frequency, a total difference of 200 ppm. In practice, the difference in nominal clock frequencies is typically much lower than 200 ppm, but since even the smallest differences in clock frequencies will cause unwanted packet loss when testing a maximum load, it is advisable to eliminate packet loss caused by small differences in line rates between the Xena tester and the DUT.

The advised approach is therefore to always reduce the line rate of the sending Xena test ports with 200 ppm, when planning for testing at 100% load. Alternatively, the test engineer can manually determine the exact ppm difference in line rates, and fine tune the line rate speed reduction to the exact level required to eliminate unwanted packet loss.

This manual fine tuning approach of the test port line rate is illustrated in the following test example, where the difference in line rates between the Xena tester and the DUT is only 4 ppm, but still enough to introduce unwanted packet loss during testing performed at 100% load.

Step 1: Select two test ports. Reserve a source and a destination port on the Xena tester. The source port is Chassis XB-2 / Module 7 / Port 3, and the destination port is Chassis XB-2 / Module 7 / Port 4. See Figure 2.

Step 2: Teach the DUT the MAC address of the destination port. On the destination test port (Chassis XB-2 / Module 7 / Port 4), setup a traffic stream sending packets to the source port, so that the DUT learns the MAC address of the destination test port. Then start the traffic generation to send the traffic from the destination test port to the source test port to teach the DUT the addresses. Unlink the Traffic Start on the destination port from the global Start button to ensure that traffic will keep running, so that the MAC learning table entries in the DUT do not expire. See Figure 3, and Figure 4.

Step 3: Define traffic on the source test port. On the source test port (Chassis XB-2 / Module 7 / Port 3), define a simple Ethernet stream with a fixed 64-byte packet length and 100% uniform traffic load. See Figure 5.

Step 4: Compare the number of transmitted and received packets. Start traffic, and compare the number of transmitted packets on the source port, against the number of received packets on the destination port.

With 64B packets, and a standard minimum Inter Frame Gap (IFG) of 20B (including the preamble), the expected number of packets per second can be calculated as $125,000,000 \text{ Mbps} / (64\text{B}+20\text{B}) = 1,488,095$ packet per second (pps). We now compare the number of transmitted packet with the number of received packets, and we can see that while the Tx rate is 1,488,095 pps as expected, the received rate is only 1,488,089 pps (see Figure 6).

We can therefore calculate that the nominal clock frequency difference between the Xena tester and the DUT as $(1,488,095 - 1,488,089) / 1,488,095 = 4 \text{ ppm}$, with the Xena tester having the highest line rate.

Step 5: Observe when sequence errors are detected. Because the Xena tester has a higher nominal line rate than the DUT, the DUT will eventually perform packet discard when the number of accumulated packets inside the DUT exceed the buffering capability. We can use this difference in line rate, to determine the buffering capability of the DUT as described in the following.

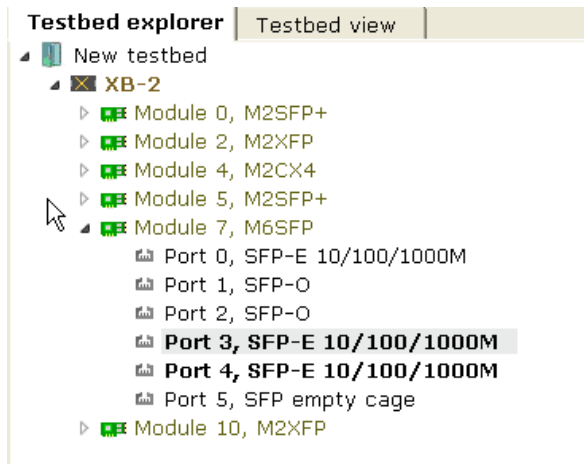
Stop the traffic generation, which will cause the buffers in the DUT to drain. Now restart the traffic, and wait for the first sequence error (=packet drop) to occur. Since the DUT accumulates 6 packets per second due to the difference in line rates, the DUT buffer size is 6 packets times the period of time from traffic generation is started, and until the first sequence error is observed. In this example, we observe the first sequence error after approximately 80 seconds. The DUT buffer size is therefore $80 \text{ sec} * 6 * 64\text{B-packets} / \text{sec} = 480 \text{ 64B-packets}$. See Figure 7.

It is important to note that this buffer size is relatively small, and we are therefore able to detect the packet loss caused by the differences in line rates after only 80 seconds of testing. If however, the DUT has provided multiple Megabytes of buffering, the period from start of testing to the first observed sequence error would have been much longer.

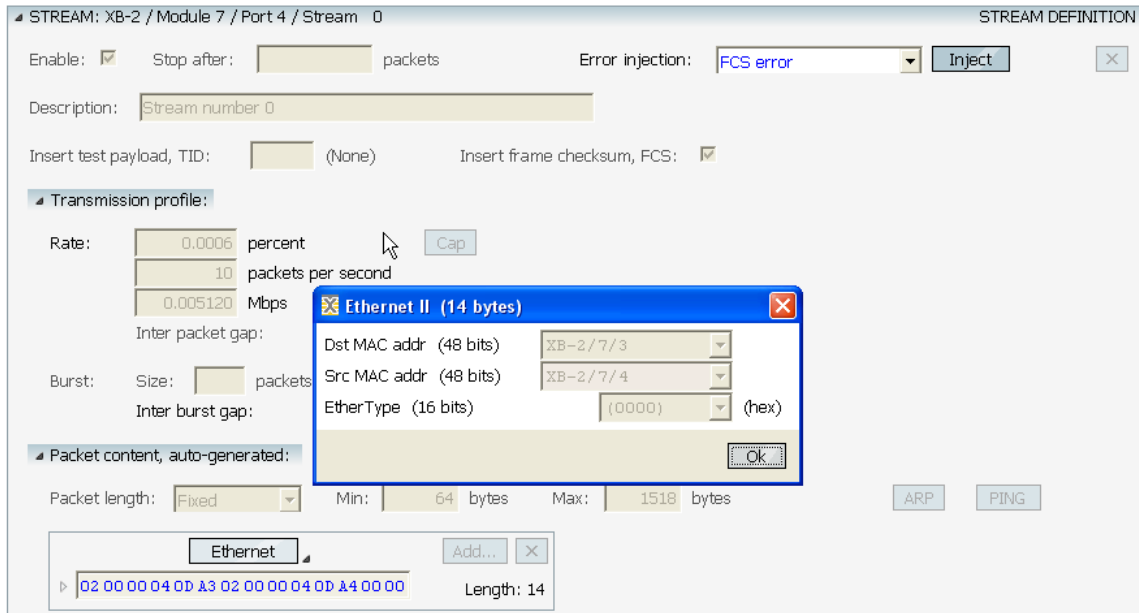
Step 6: Compensate for the line rate difference to eliminate unwanted packet loss. Stop the traffic generation, and configure the port with a 10 ppm speed reduction, which effectively reduced the line rate of the source port from 1000 Mbps to 999.990 Mbps. See Figure 8.

Since the Xena tester disables all traffic streams when the test ports line rate is changed, remember to re-enable the stream on the source port before restarting traffic. Start traffic, and let the test run for a long time to verify that the 10 ppm speed reduction of the source port can compensate for the 6 ppm difference in clock frequency between the Xena tester and the DUT. See Figure 9.


XenaManager screenshots for test steps 1 to 7




- Figure 2 Select two test ports ([step 1](#))



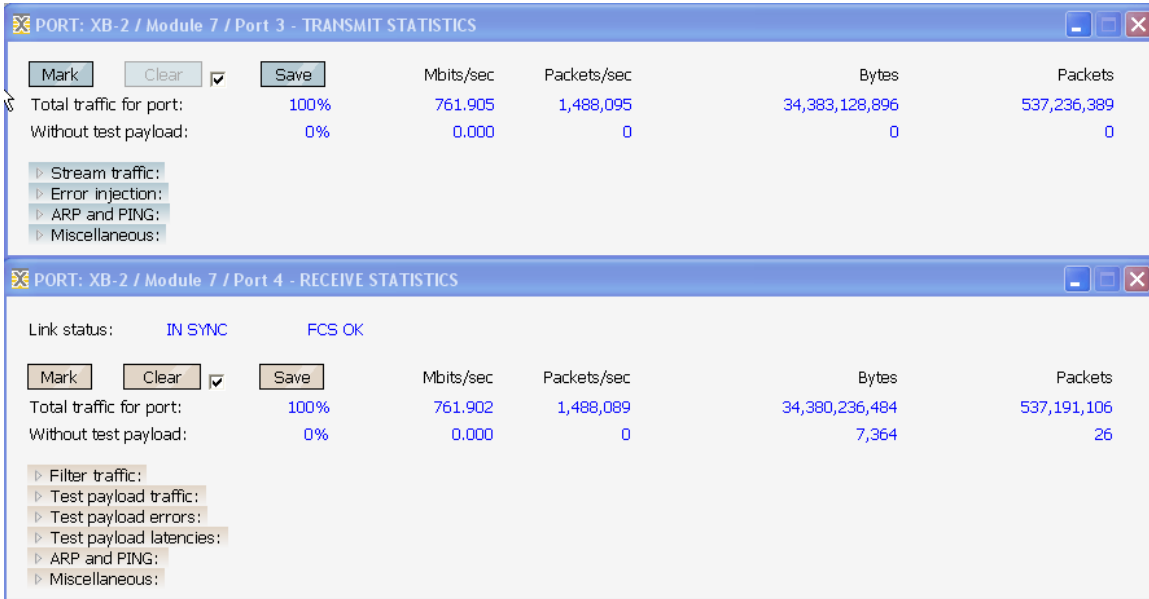
- Figure 3 Teach the DUT the MAC address of the destination test port ([step 2](#))



- Figure 4 Send the MAC address training packet forever, to avoid MAC learning table timeout ([step 2](#))



■ Figure 5 Setup traffic on the source test port (step 3)



■ Figure 6 Compare the number of transmitted and received packets (step 4)

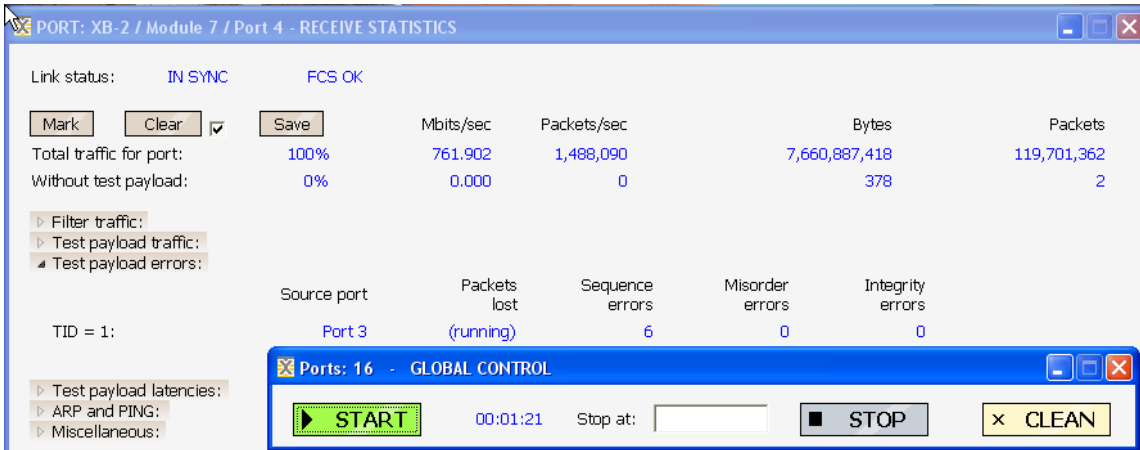


Figure 7 Observe when Sequence Errors are detected (step 5)




Figure 8 Compensate for the line rate difference to eliminate unwanted packet loss (step 6)

PORT: XB-2 / Module 7 / Port 4 RECEIVE STATISTICS

Link status: IN SYNC FCS OK

Mark Clear Save

	Mbts/sec	Packets/sec	Bytes	Packets	
Total traffic for port:	100%	761.897	1,488,080	38,890,245,316	607,659,961
Without test payload:	0%	0.000	0	9,988	34

Filter traffic:
Test payload traffic:
Test payload errors:

Source port	Packets lost	Sequence errors	Misorder errors	Integrity errors
Port 3	(running)	0	0	0

TID = 1:

Test payload latencies:
ARP and PING:
Miscellaneous:

Ports: 16 - GLOBAL CONTROL

▶ START 00:06:47 Stop at: ■ STOP ✕ CLEAN

- Figure 9 Wait for at least 80 seconds to verify that sequence error do not occur again (step 6)