New AMD Opteron™ processors and Neterion™ Xframe® Ethernet adapters first in the industry to achieve full 10 Gigabit line-rate speeds

- Opteron processors and Xframe II Ethernet adapters set a new record for network bandwidth
- Unidirectional traffic — over 9.8 Gigabits per second (Gbps)
- Bidirectional traffic — over 14.6 Gbps
- Best-in-class CPU utilization — less than 15% average per CPU on the send side

Executive Summary

This paper describes the network performance testing and results achieved with systems equipped with the AMD Opteron™ processors, AMD-8132™ HyperTransport™ PCI-X® 266 Tunnel, and Neterion Xframe® II adapters based on the PCI-X® 2.0 bus architecture with Double Data Rate technology. The AMD Performance Lab measured sustained transfer rates between back-to-back servers of about 9.86 Gbps, achieving for the first time in the industry the line-rate speed of 10 Gigabit Ethernet. Best-in-class performance for bidirectional traffic rates was also measured, as well as low CPU utilization. The server configurations and measurement methodologies are detailed, and results reported in terms of both send and receive rates for each scenario. As an additional data point, all measures obtained with Neterion's Xframe II are compared with Neterion's Xframe running on the same AMD-based systems.

Performance Results

Unidirectional Communications

	Xframe PCI-X® 1.0	**Xframe II** PCI-X 2.0 DDR
SEND SIDE (System 1: 3389)		
Total Bytes Transferred (MBytes)	107,520	107,520
Time (seconds)	113.188	87.253
Throughput (Mbits/second)	**7,599.392**	**9,858.228**
Average CPU Utilization (%)	13.67%	15.42%
RECEIVE SIDE (System 2: 3390)		
Total Bytes Transferred (MBytes)	107,520	107,520
Time (seconds)	113.640	87.832
Throughput (Mbits/second)	**7,569.166**	**9,793.242**
Average CPU Utilization (%)	47.06%	54.34%

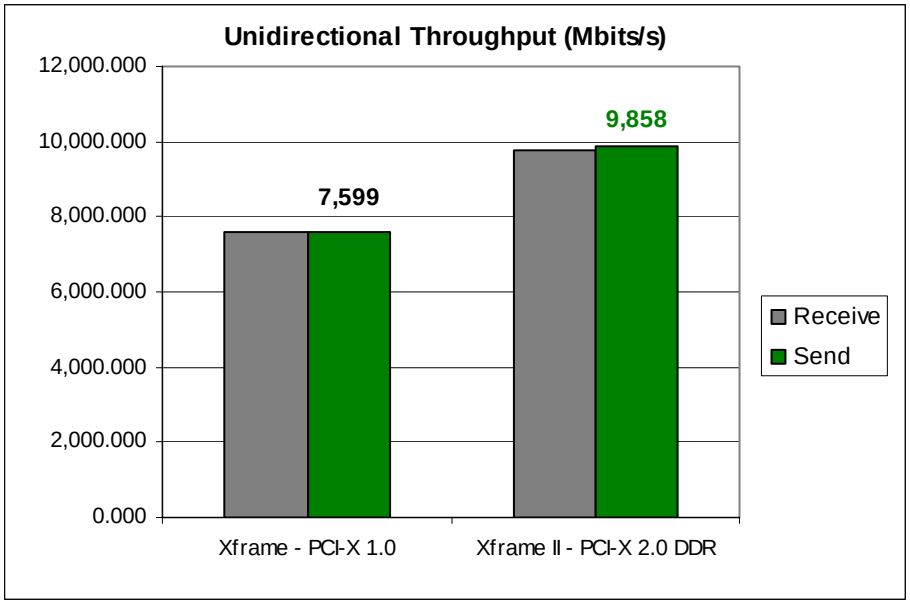

The system architecture, based on the AMD-8132™ HyperTransport™ PCI-X® 266 Tunnel, allowed for maximum allowed for maximum bus bandwidth, with throughputs close to 8 and 10 Gbps with Xframe and Xframe II respectively.

Bidirectional Communications

	Xframe PCI-X® 1.0	**Xframe II** PCI-X 2.0 DDR
Measured on System 2 (3390)		
Total Bytes Sent (MBytes)	107,516	107,520
Send Throughput (Mbits/second)	3,954.600	7,177.390
Total Bytes Received (MBytes)	107,519	107,520
Receive Throughput (Mbits/second)	3,786.612	7,461.420
Total Bytes Transferred (MBytes)	215,035	215,040
Total Concurrent Throughput (Mbits/second)	**7,741.212**	**14,638.810**

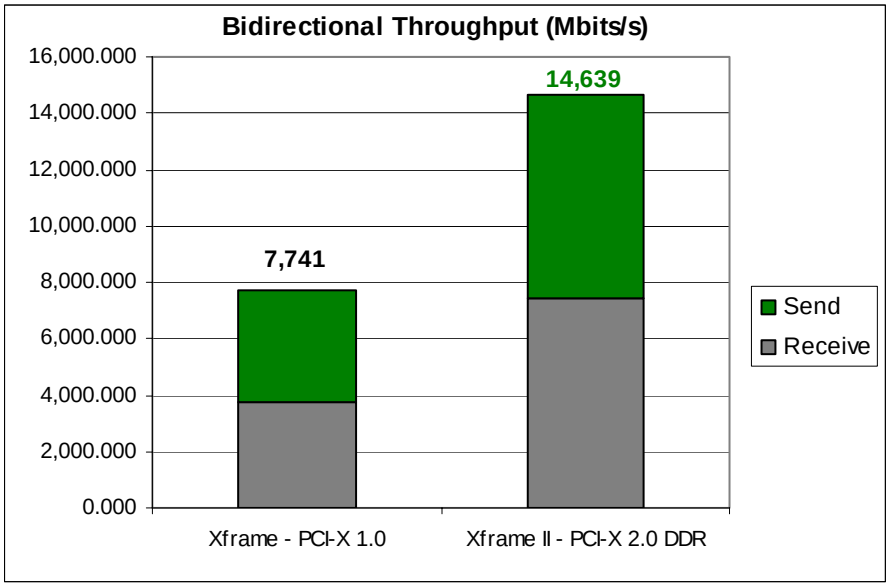

With a bidirectional throughput exceeding 14.6 Gbps, the DDR systems reached the limit of the PCI-X® 2.0 bus bandwidth.

Test Methodology

The benchmark tool used to measure the throughput is the industry-standard program: NTttcp. Available in the Microsoft® Windows® Driver Development Kit (DDK), this tool measures the transfer rate and other parameters on a TCP, UDP, or UDP multicast connection.

To optimize the transfers, the server systems were configured as follows:

- Large Send Offload (LSO): enabled (default) on both systems
- Checksum-Offload: turned on (default)
- Ethernet payload size: 9000 bytes, Jumbo frames (default: 1500 bytes)
- Resulting packet size: 9018 bytes (with Ethernet header and CRC)
- NTttcp was run on the same processor handling interrupts from the network adapter

System Configurations

Figure 1: Back-to-back setup for sending and receiving systems.

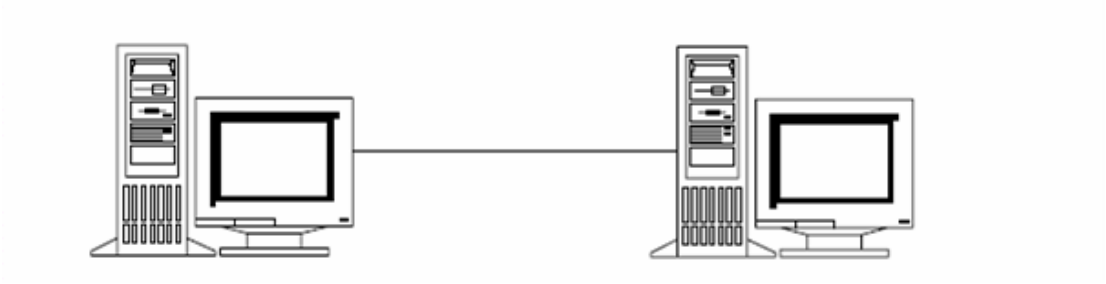

System 1 (3389):
2 AMD Opteron™ 64-bit processors at 2.8 GHz

System 2 (3390):
2 AMD Opteron 64-bit processors at 2.8 GHz

In order to eliminate external factors from the performance evaluation, a back-to-back configuration was chosen.

System Configurations

SYSTEM 1	
System	3389
CPU	(2) AMD Opteron™ Processor Model 854
Processor Clock Rate	2.8 GHz
Memory	2GB
Local Bus Specification	PCI-X® 1.0 for measuring Xframe-to-Xframe throughput; PCI-X 2.0 for Xframe II-to-Xframe II throughput measurements
Ethernet Adapters	Neterion Xframe and Xframe II 10 GbE Adapters (Xframe to Xframe; Xframe II to Xframe II)
Network Driver	Xenamd64.sys 2.0.2.0
Network Type	10 Gigabit Ethernet
Threading	Single-threaded only
Operating System	Microsoft® Windows® Server 2003, Enterprise x64 Edition [Version 5.2.3790]

SYSTEM 2	
System	3390
CPU	(2) AMD Opteron™ Processor Model 854
Processor Clock Rate	2.8 GHz
Memory	2GB
Local Bus Specification	PCI-X® 1.0 for measuring Xframe-to-Xframe throughput; PCI-X 2.0 for Xframe II-to-Xframe II throughput measurements
Ethernet Adapters	Neterion Xframe and Xframe II 10 GbE Adapters (Xframe to Xframe; Xframe II to Xframe II)
Network Driver	Xenamd64.sys 2.0.2.0
Network Type	10 Gigabit Ethernet
Threading	Single-threaded only
Operating System	Microsoft® Windows® Server 2003, Enterprise x64 Edition [Version 5.2.3790]

Performance Test Procedure

Setup

1. NTttcp installed on send and receive systems.
2. Xframe adapters installed on both sides for the first tests; Xframe II adapters installed for subsequent tests.
3. Close all applications (nothing running in the background) on both systems.
4. Registry entries added at HKLM\System\CurrentControlSet\Services\Tcpip\Parameters:
   - Registry entry Tcp1323Opts, with type as REG_DWORD, and value set to 1.
   - Add a registry entry called TcpWindowSize, also with type REG_DWORD, and value set to 512K (default 64K).
5. Systems rebooted for settings to take effect.

Use of NTttcp

Using NTttcp, throughput was measured in eight different scenarios:

- **With Xframe adapters in both send and receive systems:** NTttcp was used to measure the send and receive rates for unidirectional and bidirectional traffic.
- **With Xframe II adapters in both systems:** The same four measurement scenarios were initiated (send and receive for unidirectional; send and receive for bidirectional).

For each transfer measurement, the NTttcp tool initiated a test using Jumbo packets. The eight measurement scenarios were tested several times, without interruption, and the results averaged to obtain the numbers reported in this paper.

NTttcp Command Lines

Xframe Measurements	
Unidirectional Send Side	\nttcp_amd>ntttcps –m 1,1,1.1.1.2 -a 8 -l 537600 -n 200000
Unidirectional Receive Side	\nttcp_amd>ntttcpr –m 1,1,1.1.1.2 -a 24 -l 537600 -rb 5376000 -n 200000
Bidirectional System 3390	\nttcp_amd>ntttcps –m 1,1,1.1.1.1 -a 8 -l 537600 -n 200000 \nttcp_amd>ntttcpr –m 1,1,1.1.1.2 -a 24 -l 537600 -rb 5376000 -n 200000
Xframe II Measurements	
Unidirectional Send Side	\nttcp_amd>ntttcps –m 1,1,2.1.1.2 -a 8 -l 537600 -n 200000
Unidirectional Receive Side	\nttcp_amd>ntttcpr –m 1,1,2.1.1.2 -a 24 -l 537600 -rb 5376000 -n 200000
Bidirectional System 3390	\nttcp_amd>ntttcps –m 1,1,2.1.1.1 -a 8 -l 537600 -n 200000 \nttcp_amd>ntttcpr –m 1,1,2.1.1.2 -a 24 -l 537600 -rb 5376000 -n 200000

NTttcp Command Line Parameters

-m	Number of threads (1 for single stream), on the specified processor (0, 1, 2, or 3 on a quad system), and the local IP address for the channel being exercised.
-a	Number of buffers to be pre-posted.
-l	Length of each buffer (a multiple of the packet payload size).
-rb	TCP receive window size.
-n	Controls the duration of the test transfer (adjusted to approximate 2 minutes).

For More Information

Additional information on the NTttcp benchmark can be found on the Microsoft® TechNet Website: http://technet.microsoft.com/