

Evaluation of HP 20GbE Adapters

Evaluation report prepared under contract with HP

Introduction

For blade servers, 10Gb Ethernet (10GbE) has been available as a networking interconnect speed for several years. Today, 10GbE is widely utilized in blade servers and has become the de facto standard for blade server network interconnects. However, with the increasing processing power available in the current generation of servers, along with the expanding use of server virtualization and the growing amounts of data to be processed, many IT shops require more bandwidth than a single 10GbE connection can provide.

While the future of Ethernet includes 40Gb and 100Gb technologies, many IT shops today need more than 10Gb Ethernet but are not ready to adopt 40Gb or 100Gb Ethernet at the server. In some cases, these technologies are not readily available as supported options on their existing network infrastructure or switching gear.

In order to accommodate growing demands for bandwidth, HP has introduced its 20GbE technology for the HP BladeSystem environment. This HP 20GbE technology is compatible with today's 10GbE and 40GbE technology, and is supported in existing HP BladeSystem enclosures.

HP commissioned Demartek to analyze the performance of its 20GbE technology with its BladeSystem platform and compare it to existing 10GbE technology available today.

Key Findings

We found that the HP 20GbE technology provided twice the performance as 10GbE technology while using the same amount of physical space, and is backward compatible with existing 10GbE infrastructure. These new adapters also support SR-IOV and have full FCoE and iSCSI offload capabilities.

With list prices for the 20GbE adapters approximately 40% higher than the 10GbE adapters, the performance is doubled but not the price.

20GbE Explained

In order to develop and deploy 20GbE, HP took the existing IEEE 802.3ba specification for 40GbE and simply divided the 40Gb connection into two independent 20GbE connections. To explain how this works, we first need to provide some background on today's 10GbE.

Ethernet Today – 10Gb Per Lane

The fastest Ethernet connection in a single lane today is 10Gbps in one direction. 10GbE transmits at 10.3125 Giga-transfers/second and uses 64b/66b encoding to transmit data over the wire (copper or fiber optic).

In order to get 40GbE today, four lanes of 10GbE are bundled together to form a single 40Gb connection. For 100GbE today, a similar process is used to bundle ten lanes of 10GbE together. These technologies transmit at the same 10.3125 Gbps rate and use the same encoding scheme as 10GbE, but simply bundle multiple lanes together to achieve the desired aggregate throughput.

20GbE

HP developed 20GbE for the HP c-Class BladeSystem platforms in order to provide higher bandwidth to meet today's increasing requirements in converged and virtualized environments. This 20GbE technology is contained entirely within the HP c-Class BladeSystem components, communicating internally at 20Gb. The HP Virtual Connect fabrics connect seamlessly to external 10GbE infrastructure.

Using the standard four-lane mode used by 40GbE, HP follows the same standard but uses only the first two lanes for one 20Gb connection. The second pair of lanes is used as an independent 20Gb connection. The same advanced features provided by 40GbE are available to be used by 20GbE. HP follows IEEE 802.3ba with only minor changes to a few clauses of the specification to support 20GbE while providing interoperability with 10Gb and 40Gb link partners.

HP enables two ports of 20Gb (full duplex) per mezzanine card on its Generation 8 blade servers. The 20Gb of bandwidth in each of these ports can be used entirely for Ethernet, or can be converged with other protocols like FCoE and iSCSI. Similar to HP's Flex-10 and 10Gb FlexFabric technologies, the new adapters can be divided into four connections with dynamic, precisely assigned bandwidth allocations for each.

To obtain 20Gb performance, the new HP FlexFabric adapters must be connected to the new HP Virtual Connect Flexfabric-20/40 F8 module. This combination is compatible with any industry-standard Ethernet or Fibre Channel technology.

Speed Sensing and Investment Protection

These new HP 20GbE adapters also support multiple speeds and allow for a smooth transition to higher bandwidth, while auto-negotiating the speed with the switching infrastructure. If these new adapters are installed in a blade server connected to a 10GbE switching infrastructure then these adapters will operate at 10GbE. If the switch infrastructure is upgraded to 20GbE technology such as the HP Virtual Connect FlexFabric-20/40 F8 module, or if the blade server is moved to a different chassis that has the 20GbE infrastructure, then these adapters will automatically run at the higher 20GbE speed.

These new HP 20GbE adapters and HP Virtual Connect FlexFabric-20/40 F8 modules are supported in existing HP BladeSystem enclosures today. These newer 20GbE adapters provide investment protection by allowing the adapters and switching infrastructure to be upgraded independently.

Each port of the HP 20GbE adapter operates at the full rated speed of 20Gbps (full duplex). This is not to be confused with load balancing across two 10GbE ports, where two 10GbE ports share the load and could, in some situations, achieve 20Gbps across both ports together.

Additional information for HP Virtual Connect is available at www.hp.com/go/virtualconnect.

Optimizing MAC Addresses and Host Ports

Because a 20GbE adapter provides twice the available bandwidth of a 10GbE adapter, fewer 20GbE adapters are required to achieve a specific bandwidth level. This means that during the initial MAC address learning phase of transmissions, fewer broadcast floods are issued than would be issued with a higher number of slower 10GbE adapters.

In addition, one 20GbE adapter could replace two 10GbE adapters, freeing up precious blade server mezzanine card locations and making them available for other types of adapters, or enabling servers with additional 20GbE adapters for demanding workloads connections.

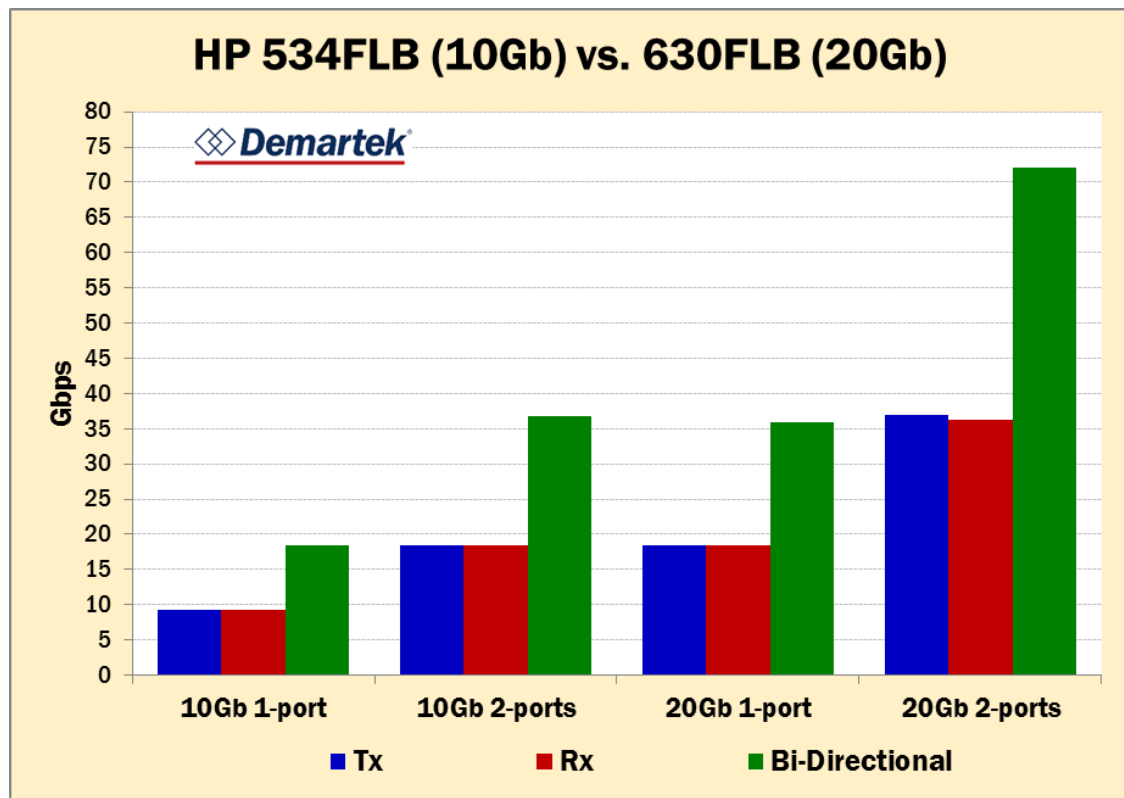
20GbE Performance Results

Four tests were performed with 10GbE and 20GbE adapters with the HP blade servers:

- ◆ 20Gb single port
- ◆ 20Gb dual port
- ◆ 10Gb single port
- ◆ 10Gb dual port

The tests were conducted using Chariot, a well-known networking test tool that sends Ethernet traffic across a network connection. Tests were run and results were recorded for transmit only (Tx), receive only (Rx) and bi-directional (simultaneous transmit and receive). Multiple sets of tests were run, and the results were very similar for each set of tests. The results of one of these sets of tests are shown below.

As expected, one port of 20Gb Ethernet performed approximately the same as 2 ports of 10Gb. Two ports of 20Gb Ethernet provided approximately double the performance of two ports of 10Gb.

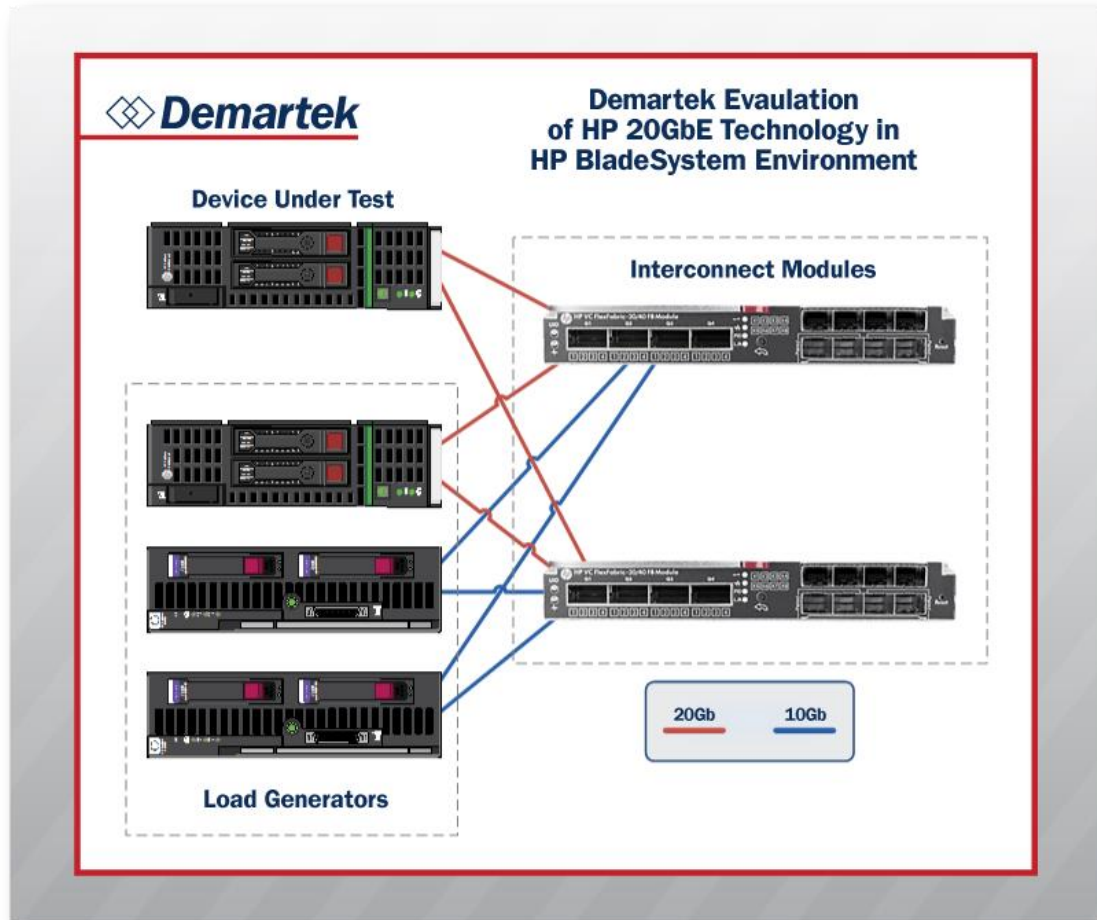


Conclusion

HP has provided an innovative way to achieve higher Ethernet network bandwidth by providing 20GbE adapters and Virtual Connect FlexFabric interconnect modules for its BladeSystem platforms. This 20GbE technology is backward compatible with today's 10GbE technology.

HP's new technology offers good value by doubling performance with only a 40% increase in price.

Appendix – Test Configuration Details



Device Under Test

HP BL460c Gen8

- ◆ 2x: Intel Xeon E5-2680 v2, 2.8 GHz, 20 total cores, 40 total logical processors
- ◆ 32GB RAM
- ◆ BIOS setting: Maximum Performance & Static High Performance
- ◆ HP 534FLB (10Gb) and HP 630FLB (20Gb) Ethernet adapters
 - ◇ HP 630FLB based on NetXtreme II network adapters from QLogic

Interconnect Modules

HP Virtual Connect FlexFabric-20/40 G8 Module

Load Generators

1x HP BL460c Gen8

- ◆ 2x: Intel Xeon E5-2680 v2, 2.8 GHz, 20 total cores, 40 total logical processors
- ◆ 32GB RAM
- ◆ BIOS setting: Maximum Performance & Static High Performance

2x HP BL460c G7

- ◆ 2x: Intel Xeon X5675, 3.06 GHz, 12 total cores, 24 total logical processors
- ◆ 48GB RAM
- ◆ BIOS setting: Maximum Performance & Static High Performance

The original version of this document is available at:

http://www.demartek.com/Demartek_HP_20GbE_Technology_BladeSystem_Evaluation_2014-06.html on the Demartek website.

Demartek is a trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.