PCI Gigabit Ethernet Performance
The Design Makes the Difference!

Includes Gigabit Performance Test Results Using
HP 9000 N-Class (Mid-range) Servers

September 15, 2000

U.S.A.

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this white paper, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Introduction

HP leads the industry in overall PCI Gigabit Ethernet (GbE) performance and reliability with the balanced system approach built into the HP 9000 N-class servers.

The goal of this paper is to show what HP customers can expect in GbE performance from N-Class servers, using HP PCI 1000BASE-T and PCI 1000BASE-SX GbE Network Interface Cards (NICs). This paper discusses the following GbE performance issues:

- **Tested Products:** Overview of the HP GbE products used to conduct the performance tests described in this article.
- **Server-to-Switch GbE Throughput Performance:** Performance data generated using an N-Class server with HP and Cisco switches, and multiple, 100BASE-TX clients. HP has achieved leadership in attaining single card link speed performance on its PCI Gigabit Ethernet products with lowest CPU utilization.
- **Back-to-Back GbE Throughput Performance:** Performance data generated using five HP N-Class servers.
- **Balanced System Design:** High speed networking cards such as Gigabit Ethernet put tremendous load on the server. Attaining optimum network performance requires a balanced system design where the I/O performance scales with the faster server processors and system bus frequencies.

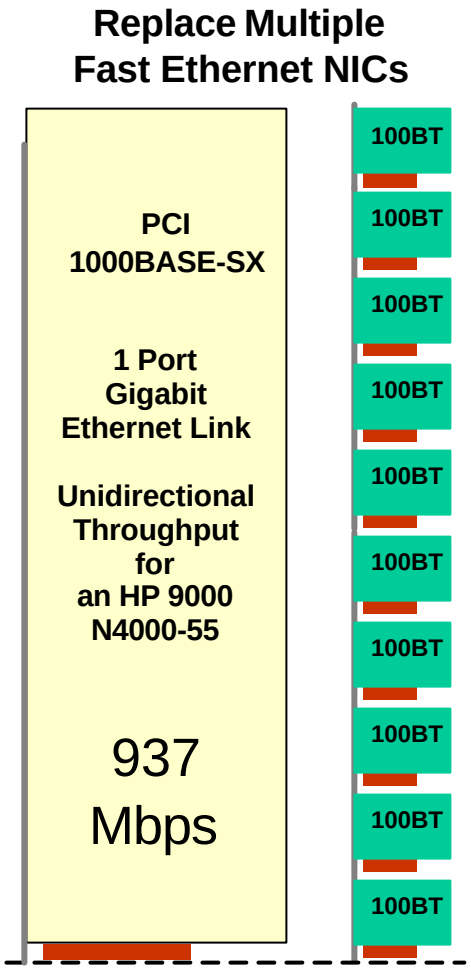

Some of the system implementations of today use PCI I/O subsystems that are hampered with respect to system I/O performance due to the use of legacy, shared PCI architecture topologies. HP's balanced system approach overcomes these legacy problems with a point-to-point, switch-like I/O topology in the HP 9000 N-Class servers.

This paper will discuss how HP servers and GbE NICs are used to offer world-class Gigabit Ethernet performance with a balanced system approach.

Tested Products

The performance test results presented in this paper were obtained with HP Gigabit Ethernet NICs and N-Class servers.

Products Used for Performance Tests With NetPerf

HP 9000 N4000-55

- 8 550MHz PA-8600 CPUs
- 8GB System Memory
- HP-UX 11.0 General Patch Release B.11.00.48

A4926A PCI 1000BASE-SX NIC

- IEEE 802.3z Compliant
- PCI-4X (64 bit/ 66 MHz)
- Jumbo Frames Support
- HP-UX 11.0 GbE Driver B.11.00.11

Note: HP also offers PCI 1000BASE-T NICs (A4929A). The performance characteristics of the 1000BASE-SX and 1000BASE-T NICs are identical.

Turbo and Twin Turbo PCI Slots

The HP 9000 N-Class servers offer two types of PCI slots:

- **Twin Turbo:** These ports support 32 and 64 bit universal PCI NICs at speeds of 33MHz and 66MHz.
- **Turbo:** These ports support 32 and 64 bit universal PCI NICs at speeds of 33MHz only.

All of the NICs used in these performance tests were HP PCI 64-bit/66MHz 1000BASE-SX and 1000BASE-T products. All were installed in Twin Turbo PCI slots.

Server-to-Switch GbE Throughput Performance

All server-to-switch tests were conducted using NetPerf and used the same HP N4000-55/8w server with A180c clients connected with one Cisco Catalyst 4000 switch, two HP 9403M routing switches, and one HP 8000M switch.

The tested system was a HP N4000-55 configured with eight 550MHz PA-8600 CPUs and 8GB of RAM. All server-to-switch links used HP PCI 1000BASE-SX NICs. All switch-to-client links used HP PCI 100BASE-TX NICs on the client side. Note: The performance obtained on the HP PCI 1000BASE-T NICs is identical to that of the PCI 1000BASE-SX NICs.

The tested system and its 48 HP 9000 A180c clients were configured with HP-UX 11.0 general patch release B.11.00.48. The HP N4000-55 was configured with HP-UX 11.0 general patch release B.11.00.48 and HP-UX 11.0 Gigabit Ethernet software driver B.11.00.11 for the Jumbo Frames and conventional 1500 byte Ethernet frame tests.

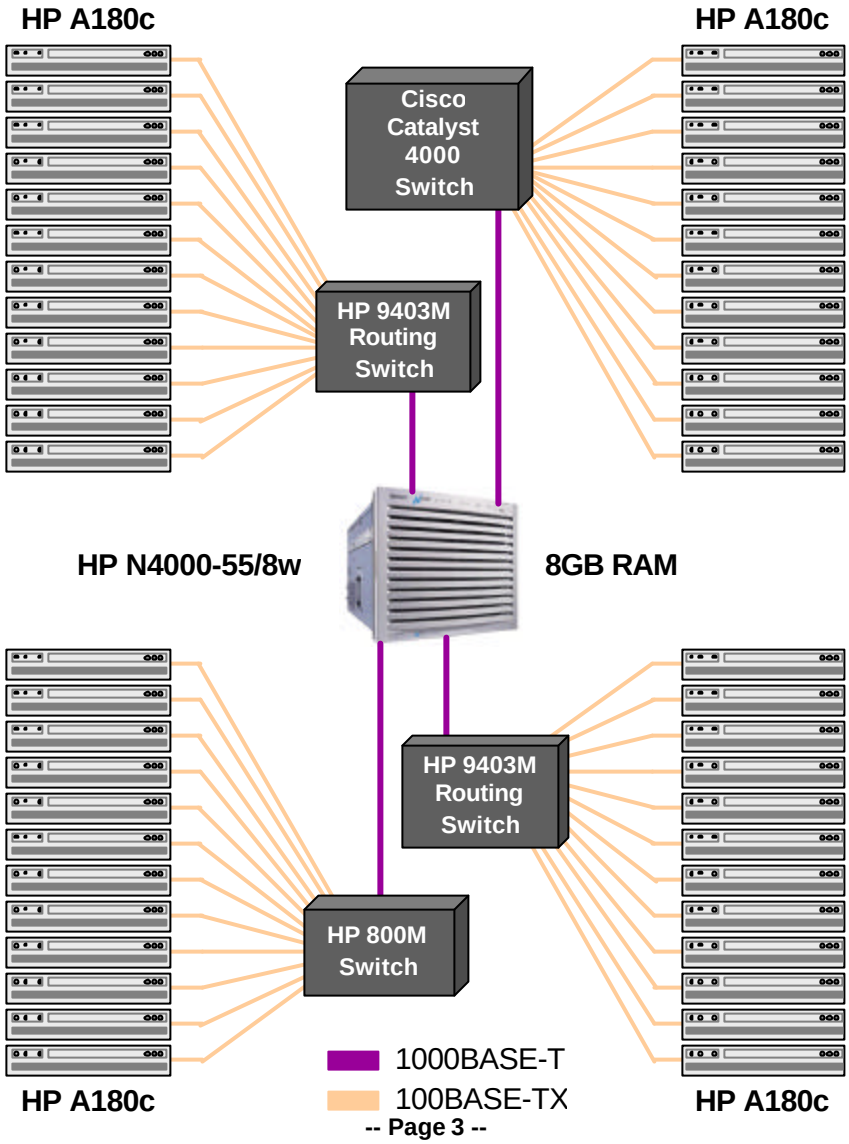

Server-to-Switch Transmit Performance

The tested system and HP A180c clients were configured with a socket size of 57,344 bytes and a message size of 16,384 bytes. All twelve HP A180c assigned to each switch were used.

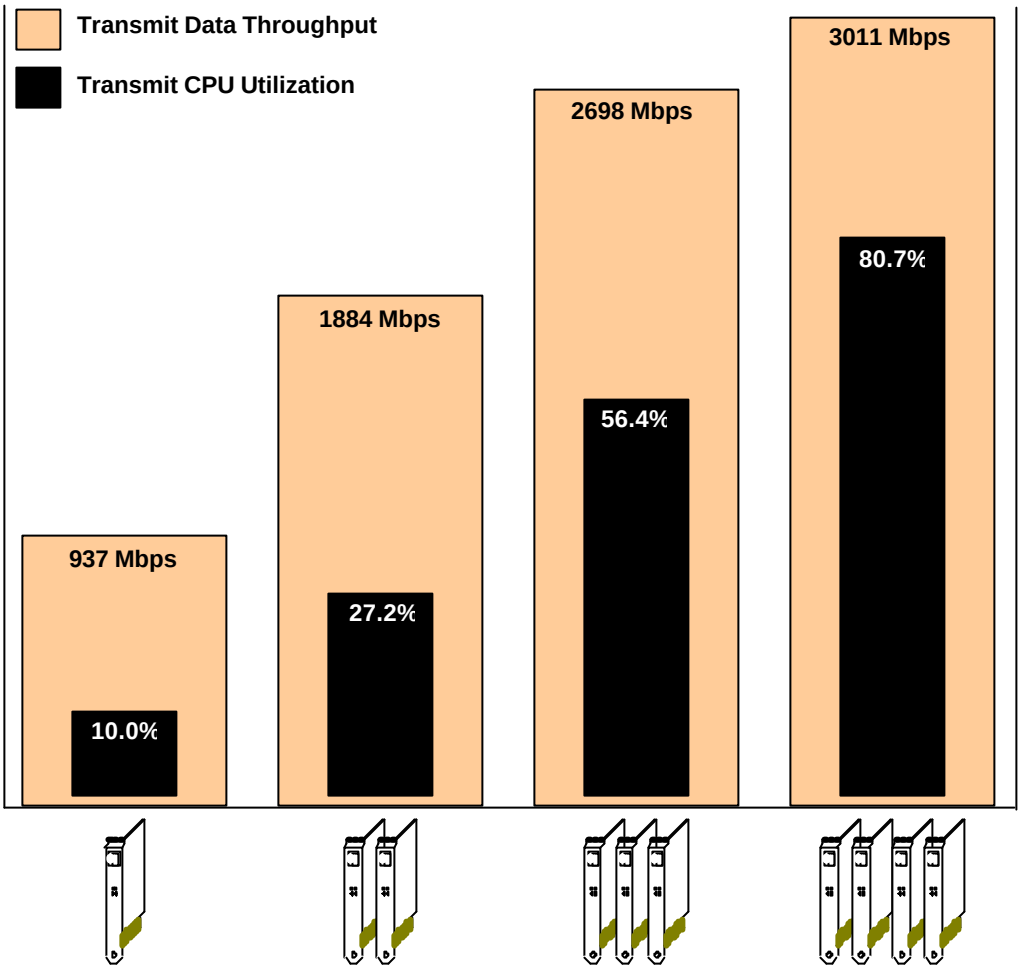

Server-to-Switch transmit throughput ratios are linear through the third NIC. Based on a factor of 1x, where 1x equals the transmit throughput of the first NIC at 937 Mbps, the throughput of the first NIC, the second, third, and fourth NICs scale at relative factors of 2.01x, 2.88x and 3.21x.

In those situations where mid-range servers like the N-Class are designated for heavy transmit roles such as media streaming and backup, the optimal number of NICs is three. Even though the fourth NIC provides additional throughput of 313 Mbps, it can be useful for failover or aggregation. In addition to providing increased throughput, this fourth NIC keeps CPU demand to a total of 80.7%.

Server-to-Switch Receive Performance

The tested system and HP A180c clients were configured with a socket size of 131,072 bytes and a message size of 16,384 bytes. Of the 12 HP A180c assigned to each switch, 10 were used.

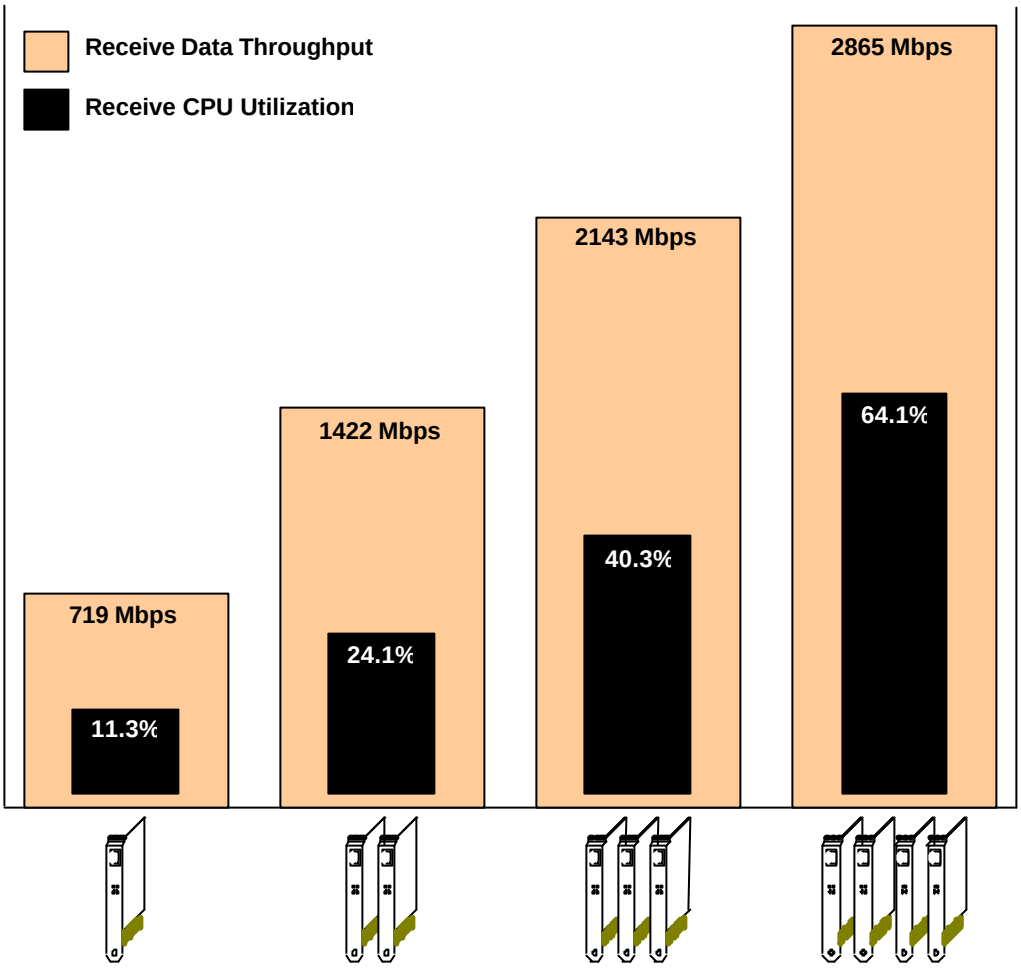

Server-to-Switch receive throughput ratios are linear through the fourth NIC. Based on a factor of 1x, where 1x equals the receive throughput of the first NIC at 719 Mbps, the throughput of the first NIC, the second, third, and fourth NICs scales at relative factors of 1.98x, 2.98x and 3.98x.

In those situations where mid-range servers like the N-Class are designated for heavy receive roles such as backup, the optimal number of NICs is four. It is important to note that an N-Class server can process 2865 Mbps of incoming data and still have 35.9% of its CPU processing cycles free for system and application tasks.

Server-to-Switch Bi-directional Performance

The tested system and HP A180c clients were configured with a socket size of 57,344 bytes and a message size of 16,384 bytes. All twelve HP A180c assigned to each switch were used.

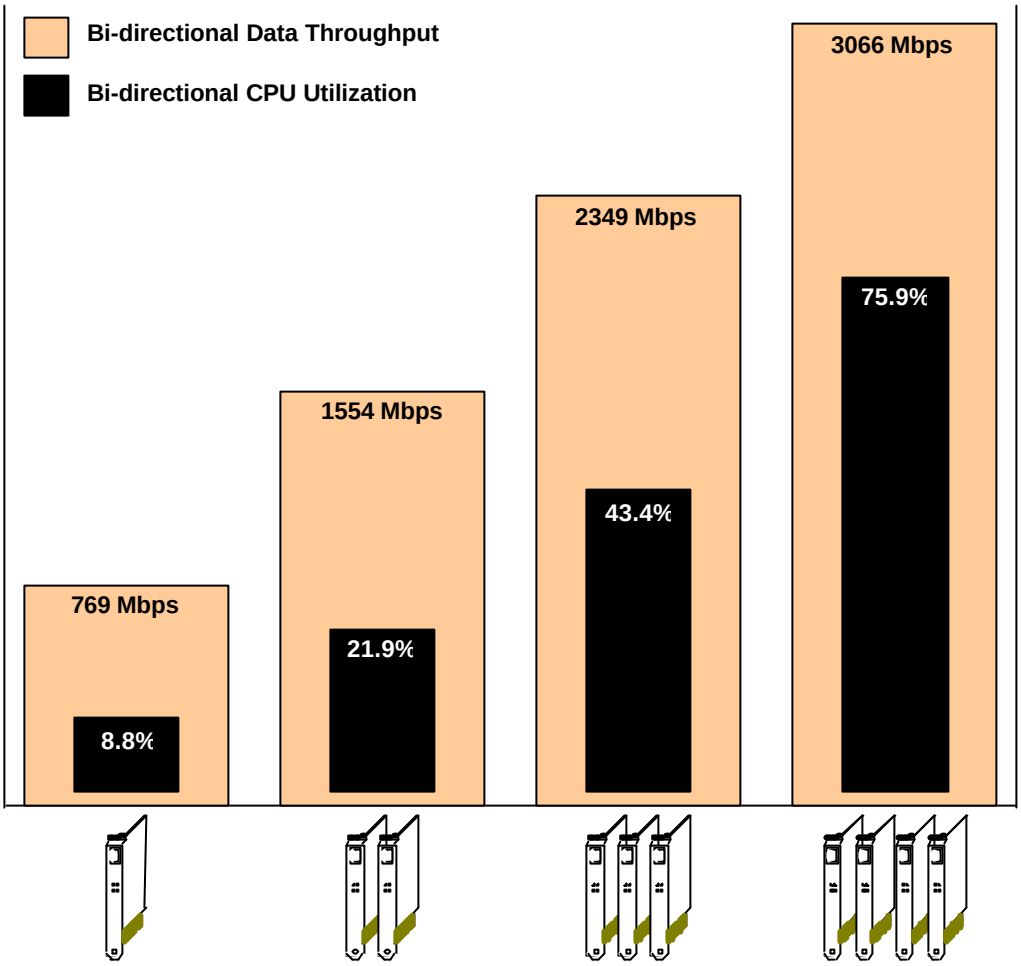

Server-to-Switch bi-directional throughput ratios are linear through the fourth NIC. Based on a factor of 1x, where 1x equals the bi-directional throughput of the first NIC at 769 Mbps, the throughput of the first NIC, the second, third, and fourth NICs scales at relative factors of 2.02x, 3.05x and 3.99x.

Bi-directional throughput performance is the “acid test” for any server, and this is where the balanced system design of the N-Class really shines with impressive linear scaling.

It is important to note that an N-Class server can process 3066 Mbps of bi-directional data and still have 24.1% of its CPU processing cycles free for system and application tasks.

Server-to-Switch NIC Throughput Scale

This throughput performance scale is based on a factor of 1x, where 1x equals the throughput of the first NIC. Scaling of the second, third, and fourth NICs is relative to the first NIC.

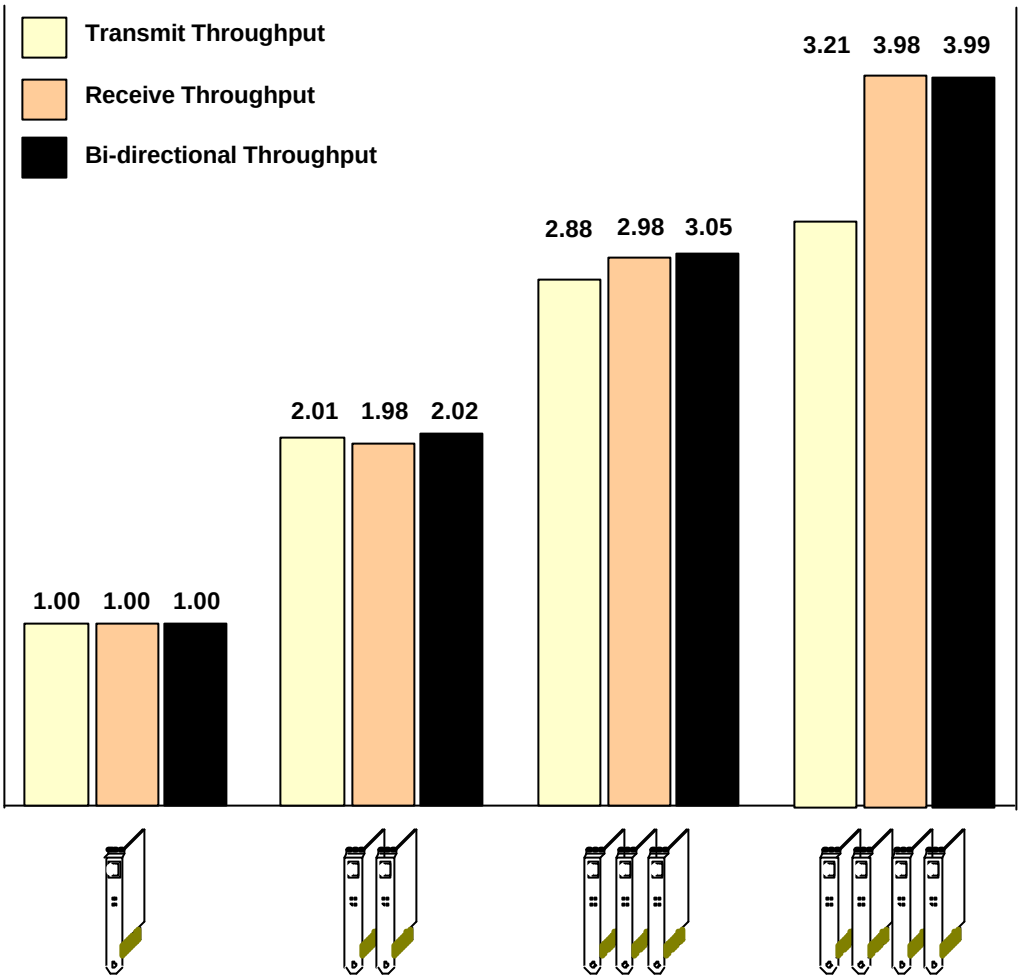

Receive and bi-directional sever-to-switch throughput performance is linear through the fourth NIC. Transmit switch throughput performance is linear through the third NIC.

As web and application servers, the N-Class servers offer solid linear performance value that is scaleable and consistently reliable because of the innovative designs of the superior PCI I/O subsystem of these versatile mid-range servers.

Back-to-Back GbE Throughput Performance

All back-to-back tests were conducted using NetPerf software, used the same three N4000-55 servers, and one HP 9000 N4000-44 server.

The tested system was an N4000-55 configured with eight 550MHz PA-8600 CPUs and 8GB of RAM system memory. All sever-to-server links used HP PCI 1000BASE-SX NICs. Note: The performance obtained on the HP PCI 1000BASE-T NICs is identical to the PCI 1000BASE-SX NICs.

The tested system and clients were configured with HP-UX 11.0 general patch release B.11.00.48 and HP-UX 11.0 Gigabit Ethernet software driver B.11.00.11 for the Jumbo Frames and conventional 1500 byte Ethernet frame tests.

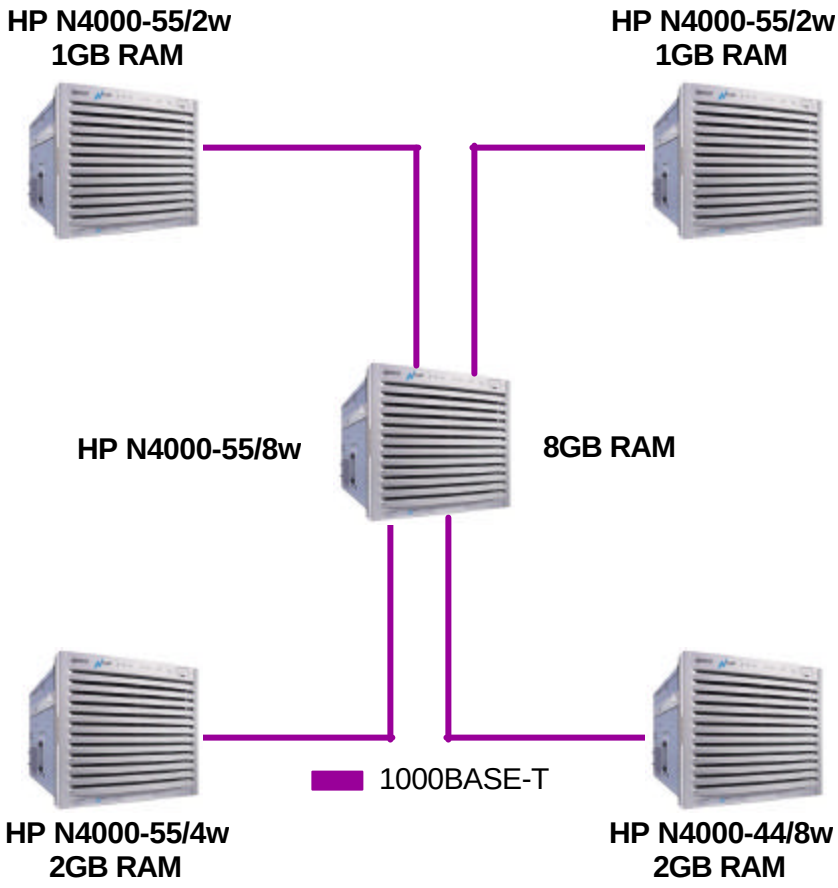

Back-to-Back Transmit Performance

The tested system and clients were configured with a socket size of 131,072 bytes and a message size of 32,768 bytes.

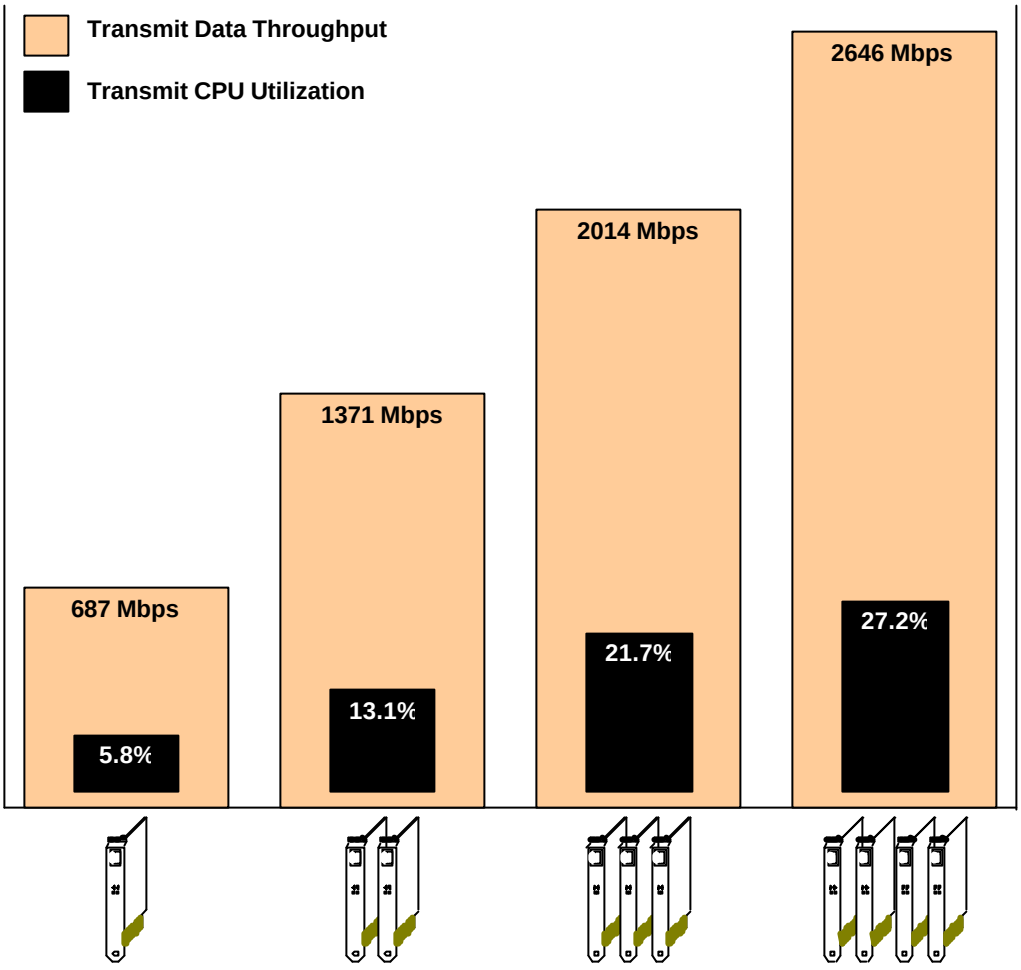

Back-to-back transmit throughput ratios are linear through the fourth NIC. Based on a factor of 1x, where 1x equals the transmit throughput of the first NIC at 687 Mbps, the throughput of the first NIC, the second, third, and fourth NICs scales at relative factors of 2.00x, 2.93x and 3.85x.

It is important to note that an N-Class server can process 2646 Mbps of transmit data over 4 Gigabit Ethernet cards in a back-to-back configuration and still have 72.8% of its CPU processing cycles free for system and application tasks.

Back-to-Back Receive Performance

The tested system and clients were configured with a socket size of 131,072 bytes and a message size of 32,768 bytes.

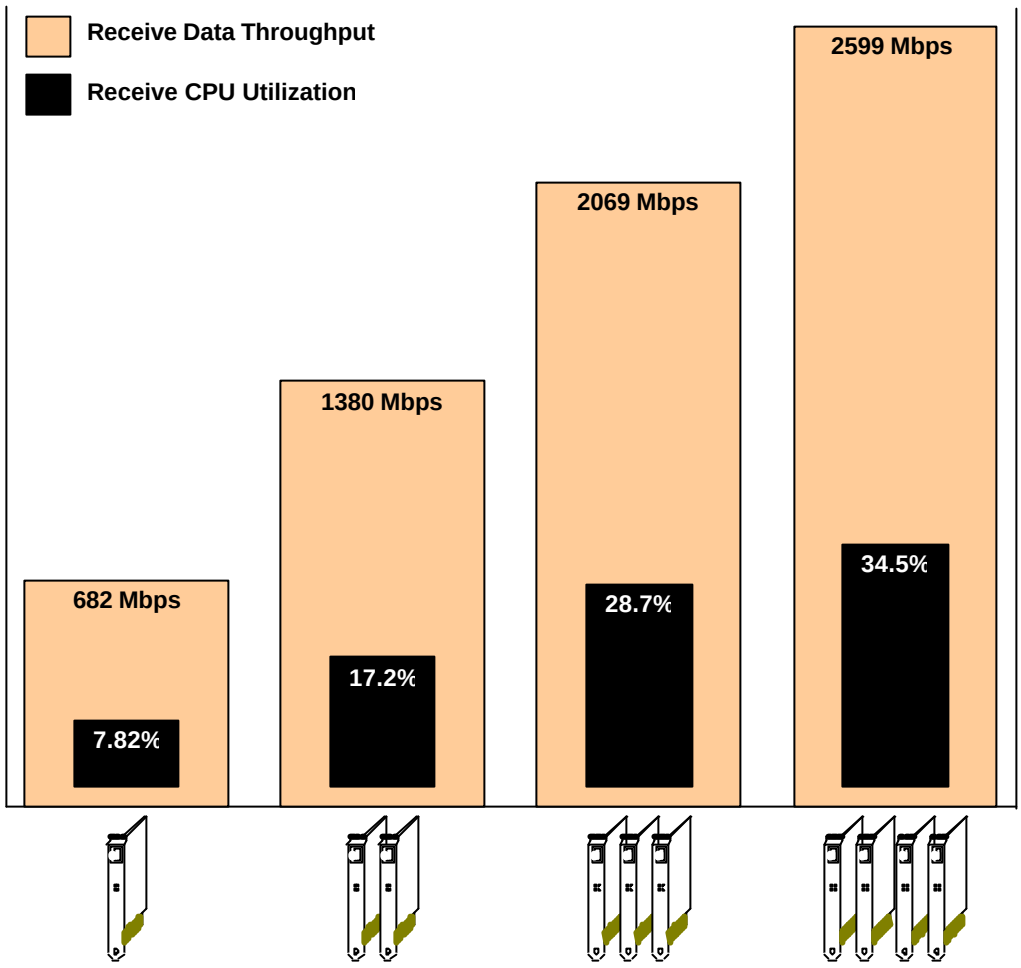

Back-to-back receive throughput ratios are linear through the fourth NIC. Based on a factor of 1x, where 1x equals the transmit throughput of the first NIC at 682 Mbps, the throughput of the first NIC, the second, third, and fourth NICs scales at relative factors of 2.02x, 3.03x and 3.81x.

It is important to note that an N-Class server can process 2599 Mbps of receive data and still have 65.5% of its CPU processing cycles free for system and application tasks.

Back-to-Back Bi-directional Performance

The tested system and clients were configured with a socket size of 131,072 bytes and a message size of 32,768 bytes.

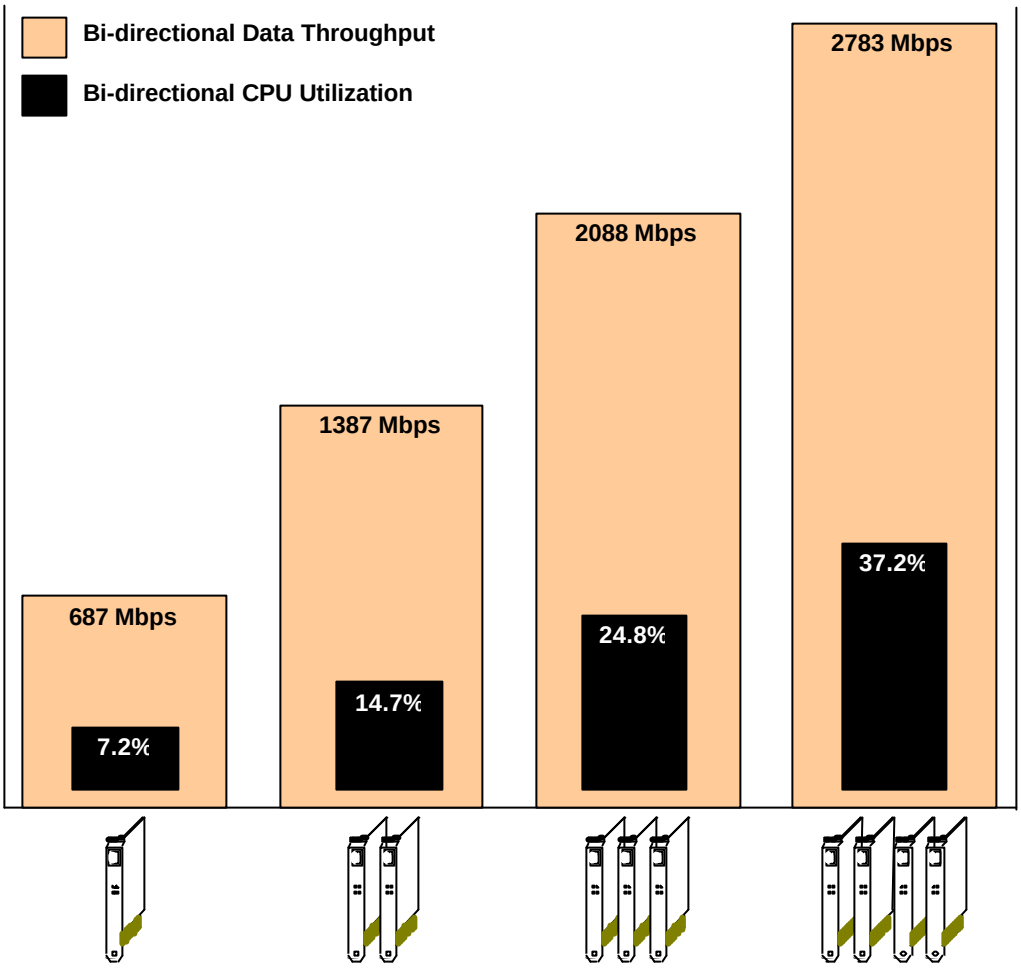

Back-to-back bi-directional throughput ratios are linear through the fourth NIC. Based on a factor of 1x, where 1x equals the bi-directional throughput of the first NIC at 687 Mbps, the throughput of the first NIC, the second, third, and fourth NICs scales at relative factors of 2.02x, 3.04x and 4.05x.

It is important to note that an N-Class server can process 2783 Mbps of receive data and still have 62.8% of its CPU processing cycles free for system and application tasks.

Back-to-Back NIC Throughput Scale

This utilization scale is based on a factor of 1x, where 1x equals the throughput of the first NIC. Scaling of the second, third, and fourth NICs is relative to the 1st NIC.

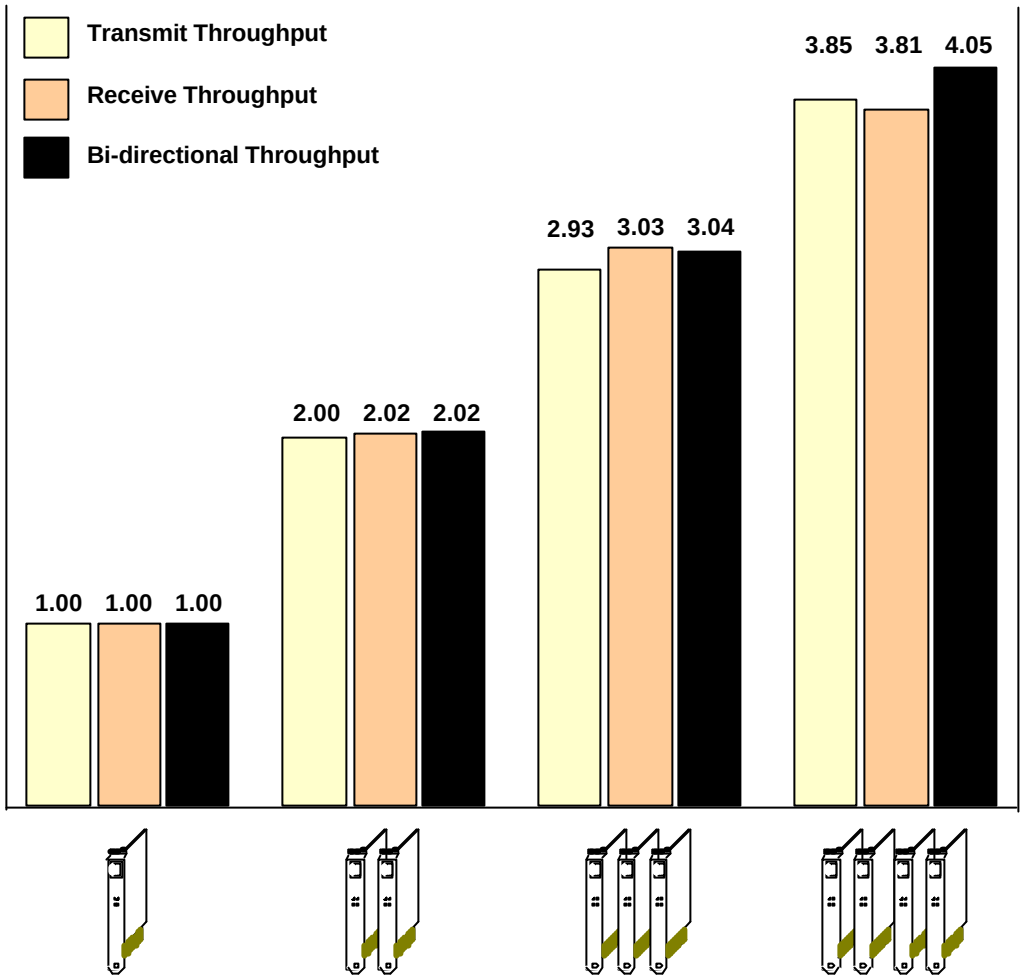

Transmit, receive and bi-directional server-to-switch throughput performance is linear through the fourth NIC with a maximum CPU utilization ratio of 37.2%. The result is that N-Class servers are back-to-back rockets because they offer scaleable, world-class throughput with a nominal CPU load.

When used in conjunction with co-located collaborating servers over PCI Gigabit Ethernet links, the N-Class servers offer solid linear performance value. This combination is scaleable and consistently reliable because of the innovative design of the superior PCI I/O subsystem of these versatile mid-range servers.

Balanced System Design

The goal of a balanced system design is to obtain comparable performance levels from the server's processor, system memory and I/O. This is exactly what every N-Class server delivers. With this design approach, the N-class servers deliver leadership Gigabit Ethernet link speed performance on a single NIC, scaleable to multiple NICs.

Optimizing PCI I/O Performance

The PCI standard was readily adopted in the early 1990s, and a major part of its widespread popularity is its ability to provide a substantial volume of I/O bandwidth. However, Ethernet speeds have increased phenomenally since the initial introduction of the PCI standard. For this reason, many legacy shared-bus designs are becoming problematic when pushed to the new high speed throughput demands of Gigabit Ethernet.

The N-Class balanced system design eliminates the annoying PCI congestion and reliability problems that continually tax some shared bus legacy designs, and replaces it with a switch-like and scaleable PCI I/O subsystem.

This innovative PCI I/O subsystem enables N-Class servers to offer world-class, scaleable Gigabit Ethernet performance as well as superior reliability, availability, and serviceability through the independent PCI I/O bus architecture.

The net result is that the HP N-Class server and PCI Gigabit Ethernet NICs can perform and scale at their maximum performance levels at all times.

A True End-to-End Solution

HP offers two different PCI Gigabit Ethernet NICs: 1000BASE-SX (up to 550m of multi-mode fiber cable) and 1000BASE-T (up to 100m of CAT5 UTP cable). Both NICs offer the same levels of throughput performance. The only difference is cost and distance.

When these NICs are used in the HP 9000 N-Class servers, the result is an end-to-end solution that offers true world-class performance.

Taking the Next Step

For a current list of tested products and a detailed analysis of your specific requirements and needs regarding HP's N-Class server and Gigabit Ethernet products, contact your HP Sales Representative.