

Boosting Server-to-Server Gigabit Throughput with Jumbo Frames



September 15, 2000

U.S.A.

©2000 Hewlett-Packard Company



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.

Copyright Notices. ©copyright 2000 Hewlett-Packard Company, all rights reserved.


Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

Trademark Notices. UNIX is a registered trademark of the Open Group.

Introduction

Conventional Ethernet offers a maximum transmission unit (MTU) of 1500 bytes. Jumbo Frames (also known as Extended Frames) offers a 9000 byte MTU. The main advantages of Jumbo Frames are:

- **Better CPU Utilization:** Jumbo Frames offers a CPU-efficient way to increase Ethernet speed and can boost the performance of heavily tasked Gigabit Ethernet servers.
- **Higher Throughput:** The 9000 byte MTU offers ATM-sized frames that are six times larger than conventional Ethernet. The larger frame sizes lead to high throughput!
- **Available Today:** All HP Gigabit Ethernet Network Interface Card (GbE NIC) products currently support Jumbo Frames without any additional costs.



This Jumbo Frames paper discusses the following:

Jumbo Frames Throughput Performance:

The increased data throughput of Jumbo Frames is dramatically demonstrated with one-NIC and two-NIC transmit, receive, and bi-directional throughput performance tests.

How Jumbo Frames Reduce CPU Overhead

Jumbo Frames reduces CPU overhead by significantly reducing the number of Ethernet packet headers required by similar-sized messages using the conventional 1500 byte MTU.

Faster Data Throughput

One entire Network File System (NFS) datagram can fit within a single Jumbo Frame, which offers the same one-to-one relationship as ATM with NFS datagrams.

Where to Deploy Jumbo Frames

Two configurations are discussed: Back-to-Back and Server-to-Switch. General interoperability considerations are also discussed.

HP Jumbo Frames is Available Now

The deployment of Jumbo Frames requires compatible network switches and server software. Jumbo Frames is already supported on the following HP Gigabit Ethernet NIC products:

- ◆ PCI 1000BASE-SX (A4926A) & 1000BASE-T (A4929A) for V/N/L/A-Class Servers and Workstations
- ◆ HSC 1000BASE-SX (A4924A & A4925A) for D/R/K-Class Servers

HP-UX 10.20 and HP-UX 11.0 support Jumbo Frames, as will all future HP-UX releases.

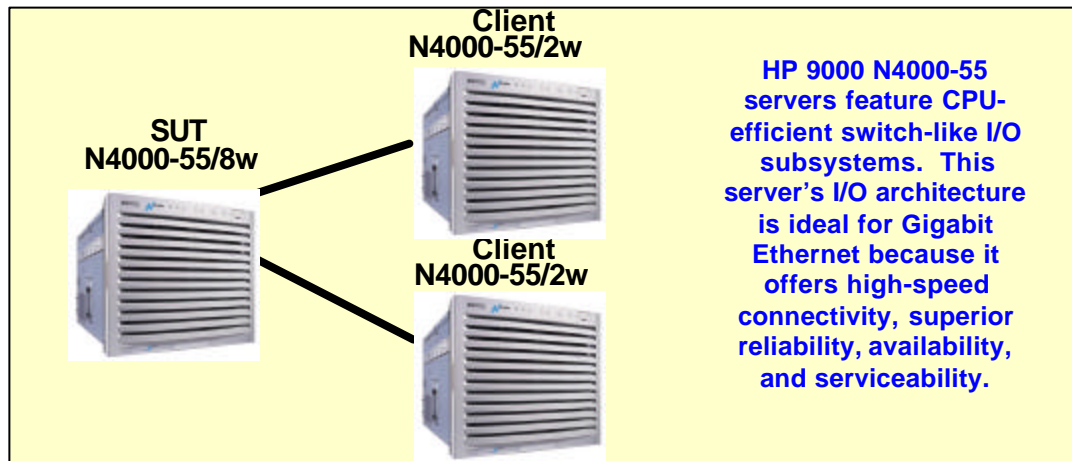
Jumbo Frames Throughput Performance

The comparison charts in this section show one-NIC and two-NIC transmit (Tx), receive (Rx), and bi-directional throughput results for Gigabit Ethernet with Jumbo Frames (9000 byte MTU) and Gigabit Ethernet with conventional frames (1500 byte MTU).

Tested Configurations

All tests used the same HP 9000 N4000-55 server with eight 550 MHz PA 8600 CPUs and 8 GB of memory as the host server. The host and remote servers 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 Ethernet frame tests.

Both PCI 1000BASE-SX and PCI 1000BASE-T NICs performed equivalently in all these tests. A back-to-back configuration was used for both the conventional Ethernet and Jumbo Frames tests. The clients are also HP 9000 N4000-55 servers with two 550 MHz PA 8600 CPUs and 1 GB of memory. All tests were conducted using NetPerf.



Jumbo Frames Configuration: Three HP 9000 N4000-55 servers (one SUT and two clients) with direct back-to-back connections and configured as follows:


- 131072 byte socket (Tx and Rx)
- 32768 byte message (Tx and Rx)

Conventional Frames Configuration: Three HP 9000 N4000-55 servers (one SUT and two clients) with direct back-to-back connections and configured as follows:

- 57344 byte socket (Tx and Rx)
- 16384 byte message (Tx and Rx)


Transmit Test Results

A server running Jumbo Frames on two NICs will transmit 603 Mbps more data, simultaneously reducing Ethernet transmit CPU usage by 44%.



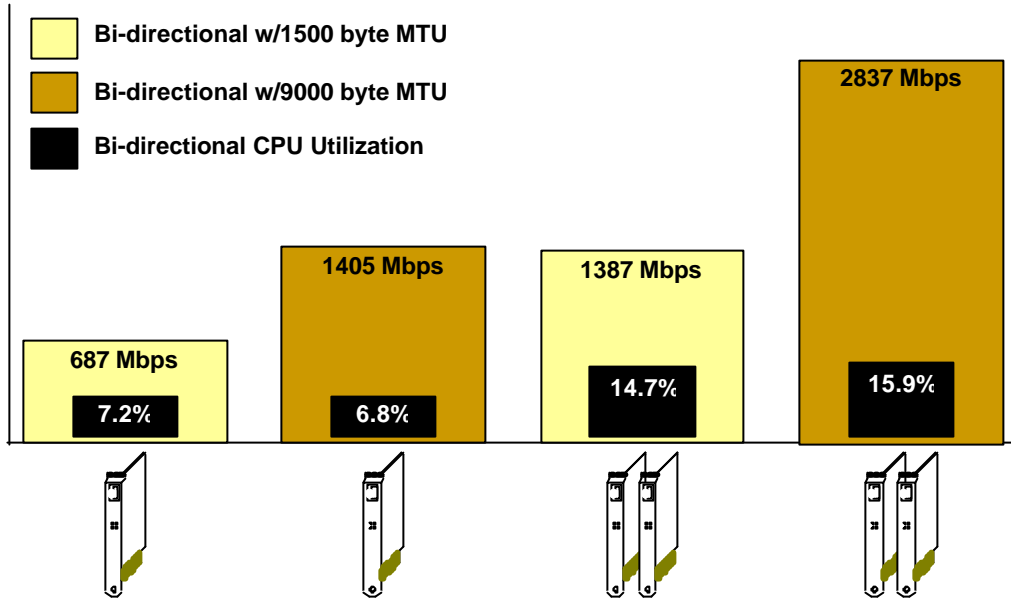
Receive Test Results

A server running Jumbo Frames on two NICs will receive 583 Mbps more data, simultaneously reducing Ethernet receive CPU usage by 27%.



Bi-directional Test Results

A server running Jumbo Frames on two NICs will send and receive (bi-directional) 1450 Mbps more data than two NICs running conventional Ethernet frame size with almost equivalent CPU utilization.



How Jumbo Frames Reduce CPU Overhead

Today's traditional and web-based servers typically handle hundreds of megabytes of data per second. However, the conventional Ethernet frame with its MTU of 1500 bytes was developed within the limitations of mid-twentieth century computers. The net result is that a substantial portion of the speed offered by today's servers is disproportionately taxed by this legacy 1500 byte MTU.

For this reason, a key benefit of Jumbo Frames is that it increases the Ethernet frame sizes to 9000 bytes. This immediately reduces the number of Ethernet packet headers that must be processed to move large files, thereby reducing the CPU overhead associated with processing packet headers.

Another way Jumbo Frames help to reduce CPU demands is by reducing the number of NIC interrupts needed when sending and receiving multi-packet file transfers.

When all of these CPU overhead reductions are taken together, the results are a gain in server performance with Jumbo Frames as well as a low cost alternative to expensive CPU upgrades in certain situations.

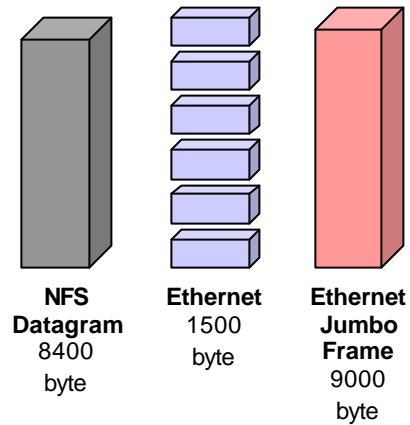
Faster Data Throughput

The ever-increasing demands of multi-terabyte data warehouses and data backups require a more efficient means of ensuring higher data rates with Network File System (NFS) datagrams.

NFS is the most widely deployed data storage mechanism used to transfer files between UNIX-based servers and features 8400 byte datagrams.

Given the 9000 byte MTU of Ethernet with Jumbo Frames, a single NFS datagram fits inside a single Jumbo Frame (1 to 1). This provides a more efficient use of memory pages with Direct Memory Access (DMA) transfers.

NFS Datagram to Ethernet Packet MTU Comparison



Comparing Jumbo Frames with FDDI and ATM

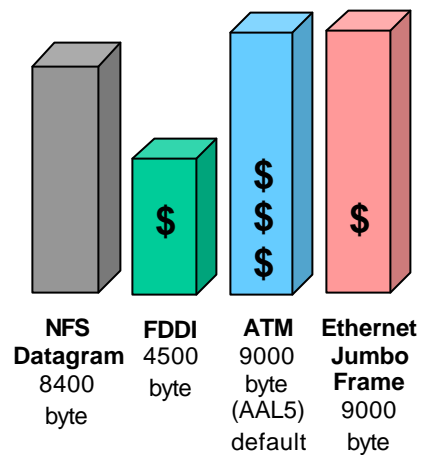
When used with Jumbo Frames, Gigabit Ethernet offers the following advantages over ATM and FDDI networking technologies:

ATM vs. Jumbo Frames

With ATM and Jumbo Frames, each NFS datagram can be sent intact. Therefore, the real difference between Jumbo Frames and ATM is cost.

Jumbo Frames Ethernet is a significantly less expensive way than ATM to transmit NFS Datagrams.

NFS Datagram Throughput and Cost Comparison



FDDI vs. Jumbo Frames

With FDDI and conventional Ethernet, the smaller frame sizes make it necessary to split the larger NFS datagrams into two or more parts, which adds CPU overhead to the transfer.

Where to Deploy Jumbo Frames

HP only recommends Jumbo Frames deployments for homogeneous Jumbo Frames topologies where all data streams pass through compatible components. This is because Jumbo Frames is still an emerging industry technology. Each product offered by each Jumbo Frames vendor must be fully tested for interoperability.

The most probable future for Jumbo Frames is that it will eventually become part of heterogeneous topologies. Until then, the data center is the best place to deploy Jumbo Frames.


Two data center configurations that have been successfully tested include:

- **Back-to-Back:** Direct Jumbo Frames Gigabit Ethernet link(s) between co-located collaborating servers.
- **Server-to-Switch:** Jumbo Frames Gigabit Ethernet links between compatible servers and switches supporting Jumbo Frames.

Some customers may choose to deploy Jumbo Frames throughout their entire organization, from the server to the desktop, to fully enjoy the cost-saving advantages of Jumbo Frames. In these types of homogeneous topologies, all data streams must pass through fully interoperable Jumbo Frames components.

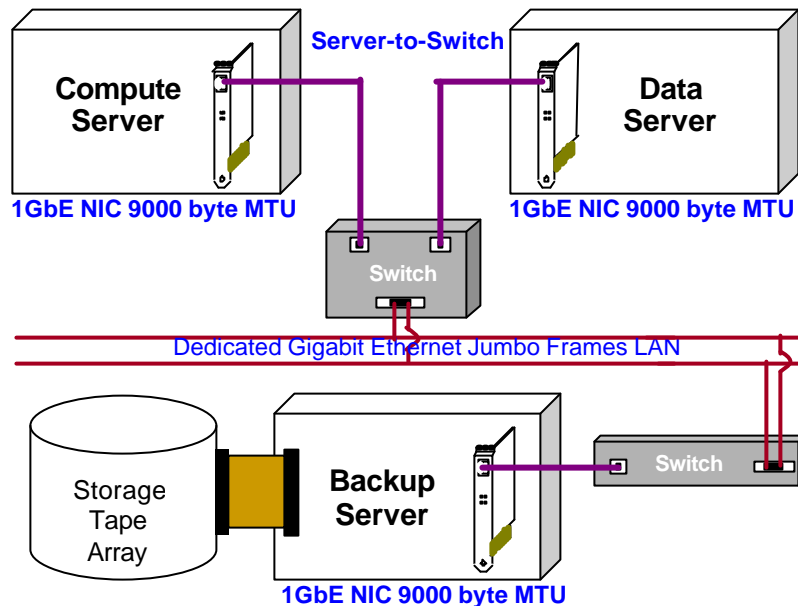
Back-to-Back Servers

All HP Gigabit Ethernet NIC products offer NIC-to-NIC links with full Jumbo Frames support. For co-located, collaborating servers, this configuration offers the quickest and least expensive way to immediately get Jumbo Frames performance within the data center.



Server-to-Switch

As switches with Jumbo Frames support become more readily available, interest will grow in the ability to link servers within a campus area using a high throughput dedicated Gigabit Ethernet Jumbo Frames LAN.



The Jumbo Frames sample configuration shown features two collaborating servers and a backup server sharing a dedicated Gigabit Ethernet LAN via switches which support Jumbo Frames.

When deploying Jumbo Frames the following rules will ensure a successful effort:


- All network components must support the Jumbo Frames 9000 byte frame size.
- All components located on the data path must support a 9000 byte MTU. *(Some components can simultaneously support 1500 and 9000 byte MTU frame sizes; but, if any component in the data path is limited to a 1500 byte MTU the data path will not support Jumbo Frames.)*
- All components must be tested for interoperability. *(HP products are fully tested and are ready for immediate use.)*

To ensure reliable customer operations, HP continually evaluates the interoperability of Jumbo Frames network products offered by other vendors.

Taking the Next Step

Jumbo Frames offers a low-cost alternative to expensive server upgrades when deployed on a 9000 byte MTU data path running at 1000 Mbps. For a data center with servers reaching saturation levels, deploying Jumbo Frames offers a quick and affordable way to gain:

- **Better CPU Utilization:** The performance increase for a heavily tasked Gigabit Ethernet server can equal that of a processor upgrade when using Jumbo Frames.
- **Higher Throughput:** Two Gigabit Ethernet NICs running a 9000 byte MTU can offer the same bi-directional throughput performance as four Gigabit NICs running a 1500 byte MTU.
- **Available Today:** All HP Gigabit Ethernet NIC products currently support Jumbo Frames. If your servers are already installed with HP GbE NICs, you can enjoy the performance gains of Jumbo Frames today, cost-free!



The attractive performance gains of Jumbo Frames cost relatively little and can economically breathe new life into heavily tasked servers. However, deploying Jumbo Frames does require a carefully designed 9000 byte MTU data path deployment plan.

For more information about Jumbo Frames and a detailed analysis of your specific requirements and needs, contact your HP Sales Representative.