Intel® PRO
Network Connections

Key Performance Considerations

For Selecting a Gigabit Server Adapter

Contents

Executive Summary

With today's explosion of network traffic combined with tight budgets, selecting the right network adapter is critical for servers. To take full advantage of a server's high-performance and high-availability capabilities, it is essential to select a server adapter, rather than adapt a desktop adapter to a server. To transmit data at faster speeds and protect existing network investments, users should consider a Gigabit Ethernet server adapter over 10/100 adapters. Based on existing Ethernet technology, Gigabit adapters feature high performance and flexibility at 10/100/1000Mbps.

Recently, network adapters have matured, and their features now include network connection fault tolerance, load balancing, and PCI hot plugs. This paper outlines criteria for selecting the right network adapter and explains several innovations in network adapters that have improved performance.

Evaluating Network Adapter Performance

A complete measure of overall adapter performance considers network throughput, CPU utilization, and host processor interrupts. To determine the best adapter for a network, it is first necessary to understand how these factors interrelate.

Network throughput and CPU utilization are the most common ways to measure an adapter's performance. Network throughput is the rate in megabits per second that packets are sent and received, while CPU utilization is the percentage of CPU capacity required to process the packets. The ratio of throughput to CPU utilization is expressed as the performance efficiency (PE) index, which was originally developed by PC Week (now eWeek) in 1995. The PE index is still the most commonly used performance ratio for rating adapters. High PE indexes, indicating high throughput with low CPU utilization, suggest favorable overall system performance.

Recently, performance labs and engineers have developed new ways to measure performance. These advancements include assessing host processor interrupts and their relation to performance.

Balancing Interrupts to Maximize Performance

Engineers have placed a new emphasis on interrupts because they affect both throughput and CPU utilization. The adapter sends interrupt signals to notify the server CPU that a packet or group of packets is ready to be processed, and to request CPU cycles for that process. A low number of interrupts per second is not necessarily a good performance indicator nor does it automatically equate to higher CPU utilization.

Too few interrupts lead to latencies that reduce throughput, while too many interrupts lead to high CPU utilization. If the adapter sends too few interrupts, packets will form bottlenecks in the receive buffer, while the system simply idles. The result is that latencies are introduced into the operation of the adapter, which in turn negatively affect server performance.

Tuning the hardware and software for optimal throughput is the key. Obviously, the rate of interrupt signals should strike a balance where it neither unduly hinders the CPU nor causes too many packets to overload the adapter.

For fine-tuning, the Intel® Adaptive Performance Tuning feature for 10/100/1000 adapters bundles an appropriate number of packets in memory before issuing an interrupt to the host (see Figure 1). The bundles dynamically adjust to the size of the traffic load, and an algorithm alters the process to the specific operating system used on the network. This reduces the number of interrupts going to the server CPU by up to 20 percent while maintaining packet data flow through the system.

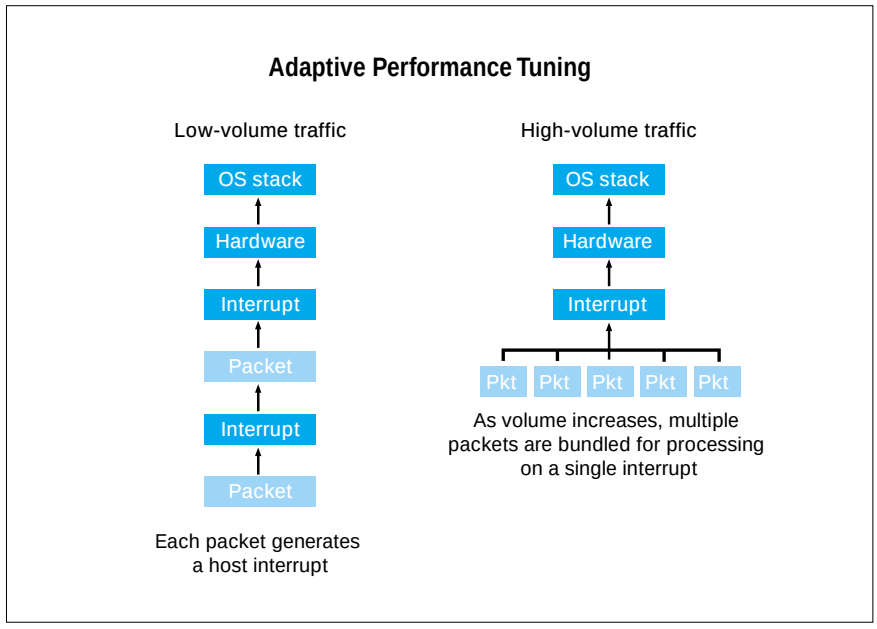

Figure 1
Adaptive Performance Tuning bundles packets to improve efficiency.

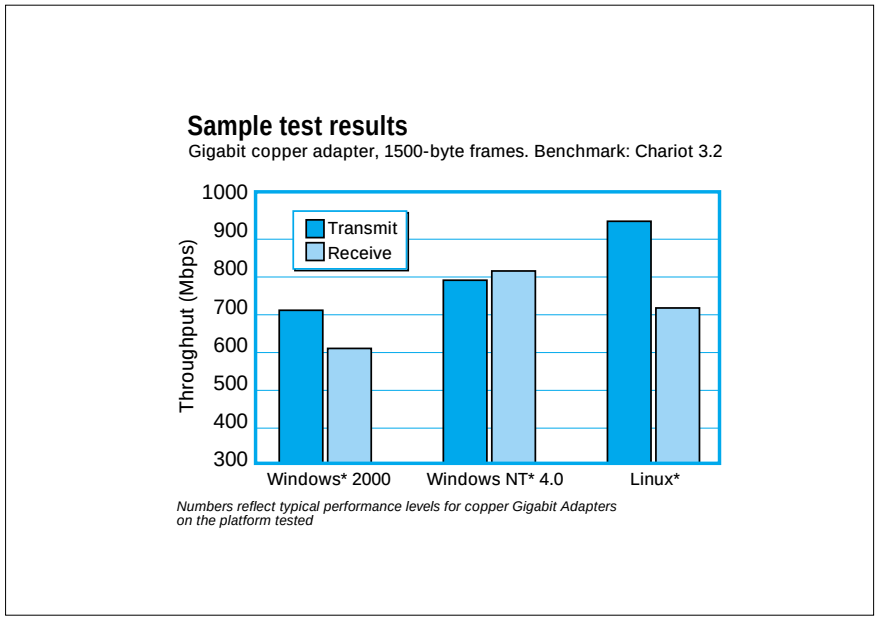

Figure 2
Sample test results for transmit- and receive-mode throughput.

Performance for Sending and Receiving Packets

Although adapters both send and receive packets, it is not uncommon for performance results to be expressed only in terms of transmit-mode throughput. Yet, to report only transmit-mode throughput is an oversight, since receive mode also requires significant work on the part of the adapter.

Since some leading applications for receive- and transmit-mode testing are interchangeable, receive-mode testing is not difficult. With popular testing applications such as NetIQ Chariot* 3.2, the same Filesendl.scr script used for transmit-mode testing can also be used for receive-mode tests, simply by reversing the endpoint pairs in the test setup. These tests typically provide receive-mode data covering both throughput and CPU utilization.

Testing receive-mode throughput is important to avoid selecting an adapter based on transmit-mode tests, only to find that it is outperformed by other offerings in receive tests. Figure 2 shows sample test results outlining the differences between transmit and receive throughput demonstrated on a Gigabit adapter with several operating systems.

Checking Operating-System Support

Perhaps the most dramatic variable for adapter performance is the operating system. This is a critical consideration because today's networks are often heterogeneous, running several operating systems in disparate locations and segments. IT managers who operate mixed network environments should check that the server adapter they choose has been tested with a variety of operating systems.

Whether in transition or planning future deployments, organizations that think strategically will better protect IT investments in the long run. For example, an organization that runs on Windows NT* today may migrate to Windows XP* a year from now to take advantage of advanced telecommunications support. Generally, it is less expensive to acquire multiple OS support in one adapter today than to add capabilities tomorrow.

A high-performance network adapter should support the following operating systems:

- Microsoft Windows* 2000
- Microsoft Windows NT* 4.0
- Microsoft Windows XP*
- Linux*
- Novell NetWare*
- UnixWare*
- Solaris*

Because many competing adapters allege multiple OS support, it is vital that adapters are tested for the full range of support. Otherwise, a user may select an adapter based on one OS, only to find that it performs poorly with an OS that may be important in the near future. It is also beneficial to look for key partnerships between switch and adapter vendors to ensure end-to-end interoperability of the hardware that provides crucial network connectivity.

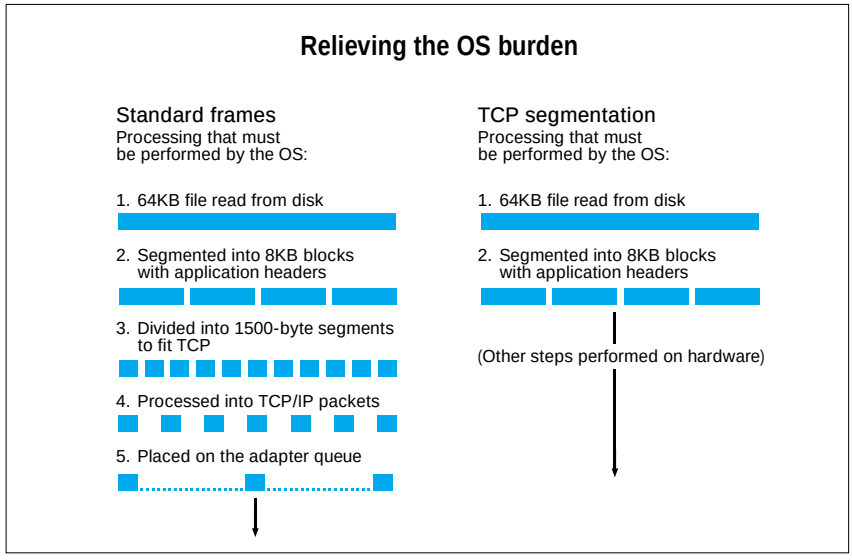

Figure 3
TCP segmentation improves performance.

TCP Segmentation Reduces Processing Loads

Along with a fully interoperable adapter, organizations should also seek adapters that include support for the new TCP segmentation or "Large Send Off-load" technology. This technology is likely to prove far more important to adapter performance than jumbo frames capability.

Both TCP segmentation and jumbo frames technologies improve performance by reducing the amount of processing that needs to be completed by the operating system (see Figure 3). When a large block of data is requested –64KB or more–the OS must break the block into packet-sized pieces, and then further process these pieces into packets that can ultimately be driven out onto the network.

With the jumbo frames feature, oversized packets are sent out on the wire, so there is less packet forming and transmitting work for the OS to do. However, with block sizes smaller than 9KB, jumbo frames performance begins to diminish as latencies are introduced.

For jumbo frames to run successfully, all nodes on the network and the switching infrastructure must be specifically enabled for jumbo frames. Because jumbo frames technology is not standards-based or ubiquitous, this feature is largely limited to niche environments.

TCP segmentation, on the other hand, does not require special end-to-end enablement because it uses regular Institute of Electrical and Electronics Engineers (IEEE*) standards-based, 1500-byte frames. It reduces the OS burden involved in processing large files into packets by off-loading this task onto hardware using specialized silicon on the adapter.

This capability is included in the new Microsoft Windows XP* operating system and is supported by Intel's latest offering of desktop and server adapters, including Gigabit adapters.
It can yield better results than jumbo frames, including higher throughput and lower CPU utilization without the necessity of configuring the entire network.

Conclusion – Selecting the Right Adapter

By looking carefully at the type and scope of testing performed, and asking for additional test information where needed, organizations become empowered to choose the best Gigabit adapter for their needs. And that, in turn, enables organizations to obtain the best performance for their server investment.

A checklist for evaluating adapter performance tests

To compare server adapters thoroughly, look for testing that covers these key items:

- Throughput (measuring both transmit and receive modes)
- CPU utilization
- Performance efficiency
- Host interrupts
- Multiple operating system support
- TCP segmentation support
- Advanced features such as Link Aggregation and virtual LANs (VLANs)

For More Information

For more information about the Intel® Gigabit Solutions, please visit http://www.intel.com/network/connectivity/solutions/gigabit.htm

For specific product information, visit the following Web sites:

Intel® PRO Server Adapters
http://www.intel.com/network/connectivity/products/server_adapters.htm

Intel® PRO/1000 XT Server Adapter
http://www.intel.com/network/connectivity/products/pro1000xt.htm

Intel® PRO/1000 XF Server Adapter
http://www.intel.com/network/connectivity/products/pro1000xf_server_adapter.htm

0302/OC/EW/PDF