

## Keeping Pace with the Rapid Adoption of Linux

With corporate IT spending levels down, IT managers need to get maximum benefit from server purchases. Increasingly, servers running on the Linux® operating system may provide the best return on investment (ROI). This article discusses how Intel and Dell are keeping pace with the increasing popularity of Linux by offering hardware and software that help enable high performance.

BY GARY GUMANOW

To obtain the maximum benefit from IT purchases, IT departments strive to invest in those projects that provide the biggest return on investment (ROI) to the organization. Now that the Linux® operating system (OS) is widely recognized as a valid choice for enterprise computing, IT managers may achieve better ROI by using systems that can run either Microsoft® or Linux operating systems. Migrating from proprietary operating systems to Linux can help firms reduce software licensing and hardware expenses.

### Forecasting dramatic growth of installed base for Linux

IDC projects the installed base of Linux to increase dramatically, and in particular, that the five-year compound annual growth rate (CAGR) for the Linux server operating environment installed base will be 24.7 percent (see Figure 1).<sup>1</sup> This rate significantly outpaces the installed base growth estimate for the other operating systems that are prevalent today.

Research from Gartner Dataquest shows where Linux is being implemented today (see Figure 2).<sup>2</sup> This data

reflects how firms are using Linux in a number of ways—ranging from Web servers to file-and-print servers.

Intel has continuously provided driver support for various operating systems, including the Linux, Novell® NetWare®, FreeBSD™, UNIX®, and Microsoft operating systems, to better meet individual business needs. As the first network adapter hardware vendor to release open source device drivers for both Fast Ethernet and Gigabit Ethernet<sup>3</sup> product lines, Intel has expanded its commitment to support Linux. Intel has developed and contributed to Linux stack enhancements that greatly enhance the performance of the OS. In addition, Intel has provided teaming technologies such as Intel® Advanced Network Services (ANS), which allow for network fault tolerance, adapter load balancing, and link aggregation. This article highlights contributions that Intel has made to improve the performance and manageability of Linux environments.

### Supporting Linux with TCP Segmentation Offload

IT managers considering Linux can benefit from choosing a server with a network connection that includes support

<sup>1</sup> Source: IDC, 2002.

<sup>2</sup> Source: Gartner Dataquest, September 2002.

<sup>3</sup> Gigabit Ethernet indicates compliance with IEEE 802.3ab and does not connote speeds of 1 Gbps.

	2001 installed base (millions)	2006 installed base (millions)	2001–2006 CAGR (%)
Combined Linux server operating environments (paid copies)	2.2	6.6	24.7%
Combined Windows server operating environments	7.0	14.3	15.2%
Other server operating environments	7.4	4.0	-11.6%

Figure 1. Worldwide server operating environments installed base, 2001 and 2006. Source: IDC, 2002

for the new TCP Segmentation Offload (TSO) technology (also known as Large Send Offload [LSO] in Microsoft operating systems). TSO, originally designed for the Microsoft Windows® 2000 OS, is now also available for Windows XP and Windows .NET Server 2003. The performance gains offered by TSO were so substantial in the Microsoft operating systems that Intel took advantage of TSO in the Linux environment. Today, the Intel 82545EM, 82546EB, and 82540EM Gigabit Ethernet controllers support TSO. Intel continues to work closely with Linux kernel developers to make TSO a robust Linux offering.

Jumbo frames (frames larger than the standard 1518 bytes) is another technology that improves performance by reducing the amount of processing that the operating system must complete (see Figure 3). When a large block of data is requested—64 KB 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.

When Jumbo frames is enabled on the network, packets going out on the network are simply oversized and the OS has fewer packets to create and transmit. However, when block sizes are smaller,

	3Q 00	2Q 01	3Q 01
Application	4	5	9
Communications	5	5	34
Database only	6	4	4
E-mail or messaging	8	10	7
External (Internet) Web	19	27	36
File or print	4	7	3
Infrastructure	22	6	12
Internal (intranet) Web	18	14	8
Multifunction	0	4	14
Others	5	7	0
Total*	8	9	12

\*Note: Total is percentage of all servers running Linux.

Figure 2. New Linux server installations according to function. Source: Gartner Dataquest, September 2002

To deploy, update, and maintain their networks, IT managers need standards-based, heterogeneous tools that they can implement across multiple platforms and operating systems.

Jumbo frames performance begins to diminish as latencies are introduced. Both TSO and Jumbo frames achieve performance gains; however, with Jumbo frames, every node on the network and all of the switching infrastructure in between must be specifically enabled.

TSO, a feature found in the OS, improves performance by off-loading packet formation from the host processor to the Ethernet network controller. TSO does not require special end-to-end enablement because it uses IEEE®

standards-based 1500-byte frames. TSO allows the host to send 64 KB blocks of data to the network adapter, instead of 1518-byte packets. The network adapter’s controller parses the 64 KB block into standard Ethernet packets, thus reducing the host processor utilization. By increasing network throughput and reducing CPU utilization, TSO increases network efficiencies.

The network adapter silicon is designed specifically for this task. Enabling TSO in Linux and using Intel Gigabit Ethernet silicon designed for TSO enhances system performance. In some cases, Intel has found a 60 percent reduction in CPU utilization and an increase in throughput. Figure 4 shows the benefit of TSO in an environment using Intel PRO network connections and Dell™ PowerEdge™ servers running Linux.

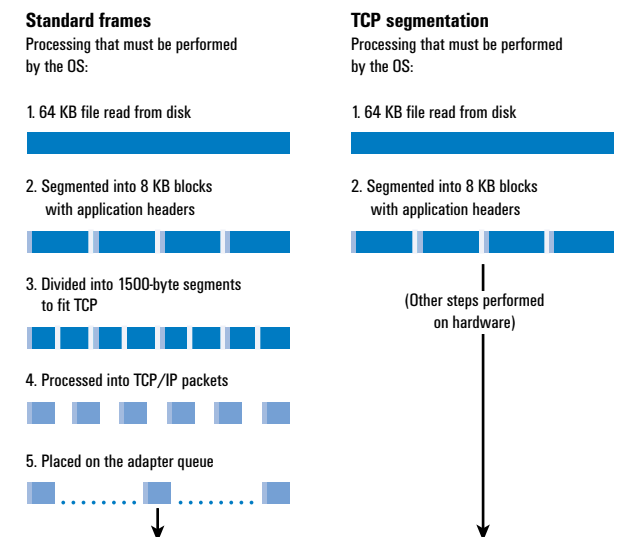


Figure 3. Relieving the OS burden: Standard frames compared to TCP segmentation

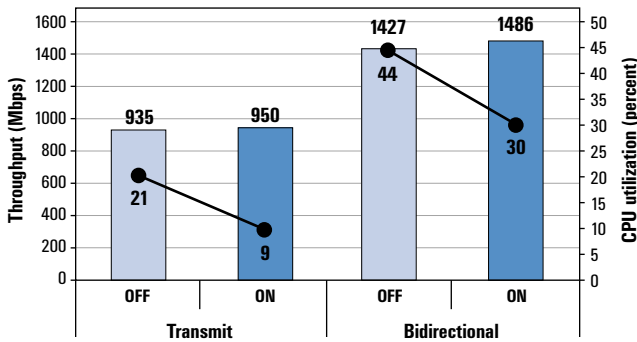


Figure 4. NetIQ™ Chariot™ network testing software results: TCP segmentation on servers running Linux, on versus off. Source: Intel Labs, September 2002

### Easing deployment with configuration tools

As more and more users begin to migrate to Linux, developers have created tools to simplify installation and configuration. Intel has contributed to such efforts through the Intel PROSet utility for Linux, a graphical user interface (GUI) that works on all Intel PRO/100 and PRO/1000 network connections (see Figure 5). PROSet for Linux displays real-time link and activity information, allows for configuration of adapter speed and duplex, performs diagnostics on the network connection, displays the Media Access Control (MAC) and TCP/IP addresses, enables installation and configuration of VLANs, and facilitates teaming of network connections. PROSet also has a feature to “blink” the LEDs on the physical adapter, aiding identification of the adapter being configured.

### Supporting Linux with device drivers

Since 1999, Intel has been developing stable and reliable high-performance General Public License (GPL)<sup>4</sup> device drivers for Linux to be used with the Intel PRO/100 and PRO/1000 network adapters and LAN on Motherboard (LOM) solutions. The Intel drivers are included with the Linux kernel, making installation of Linux easier. Intel supports these drivers by releasing frequent updates to respond to user feedback and to maintain compatibility with the latest Linux kernel patches.

### Supporting Linux with ethtool

Intel networking products include an interface for ethtool, a Linux net driver diagnostic and tuning tool compatible with the Linux 2.4.x (or later) series of kernels. The ethtool utility obtains information and diagnostics related to media, link status, driver version, and Peripheral Component Interconnect (PCI) or other bus locations.

### Advancing RASM capabilities

When purchasing a server, IT managers are likely to specify their application first and then select an OS to host the application

and to provide tools for the reliability, availability, scalability, and manageability (RASM) demands of enterprise computing. Dell offers Linux as an option on its PowerEdge product line. Working together, Dell and Intel help enterprises increase the RASM capabilities of server hardware by offering systems that include Linux drivers embedded in the kernel and systems management tools, such as Dell OpenManage™ software.

### Enabling reliability

Hot Plug PCI technology is an important step forward in creating environments that offer redundancy and fault tolerance, and it is supported on all Intel PRO server adapters. Technologies such as Intel Adapter Fault Tolerance (AFT), Intel Adaptive Load Balancing (ALB), Cisco® Fast EtherChannel®, and Cisco Gigabit EtherChannel can be used on servers equipped with Hot Plug PCI to allow replacement of failed adapters while users continue to work, thereby increasing server reliability. Hot Plug PCI is available on selected Dell PowerEdge servers.

### Enabling availability

CD-based and browser-based redeployment tools can increase availability by providing a quick recovery mechanism to restore the server to its original state. Intel AFT technology also helps increase availability by providing backup network links if a network connection fails because of factors such as loose cables or switch port failures. AFT monitors the server connection to the network (see Figure 6) and automatically switches traffic to a redundant link if a failure occurs. Administrators can create this redundancy by deploying both Gigabit Ethernet network connections found on Dell PowerEdge servers and using one as a redundant network backup link. Or, administrators can use a Gigabit Ethernet LOM that has an Intel PRO server adapter as a redundant network backup link.

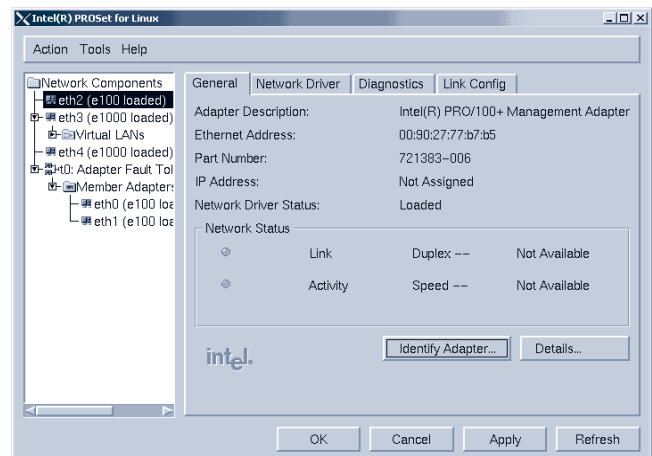


Figure 5. Intel PROSet for Linux GUI

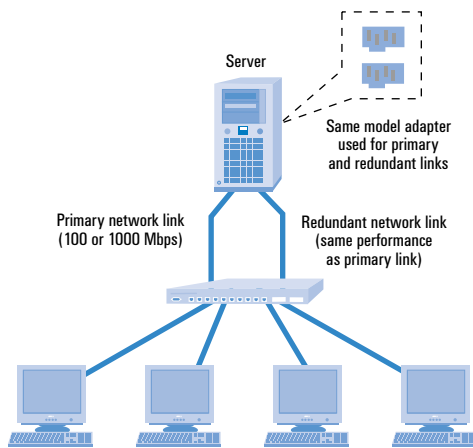
<sup>4</sup>Required for inclusion with Linux kernel releases.

### Enabling scalability


Intel has developed Adaptive Load Balancing to move data between the server and the network on demand. ALB can increase server bandwidth, up to a theoretical 8 Gbps traveling over Gigabit Ethernet, by automatically balancing traffic across as many as eight network connections (see Figure 7).

Each additional adapter adds another 100 Mbps or 1000 Mbps link to the network. The existing IP address of the server is shared by all of the adapters in the server, and traffic is always balanced between them. All the adapters in a team must be connected to one or several switches, provided the switches are on the same network segment. No client configuration is required. Once ALB is configured, all outgoing server traffic is balanced across the teamed adapters.

Intel Link Aggregation and Cisco Fast EtherChannel and Gigabit EtherChannel are other technologies that can be used to scale server bandwidth. Like ALB, these technologies automatically balance traffic and require a switch that will support Fast EtherChannel,



**Scenario 1: Full-bandwidth redundant link**



**Scenario 2: Lifeline redundant link (mixed adapter teaming)**

Figure 6. Creating fault tolerance with Intel AFT

Gigabit EtherChannel, or both. Both incoming and outgoing server traffic is balanced and can be scaled in increments. A newer industry standard for link aggregation, IEEE 802.3ad, allows for balancing traffic among multiple switches from a single server (Figure 8). All of the scalable bandwidth features supported by Intel PRO server adapters integrate seamlessly into Microsoft Windows NT®, Windows, Linux, and Novell NetWare operating systems.

### Enabling manageability

To deploy, update, and maintain their networks, IT managers need standards-based, heterogeneous tools that they can implement

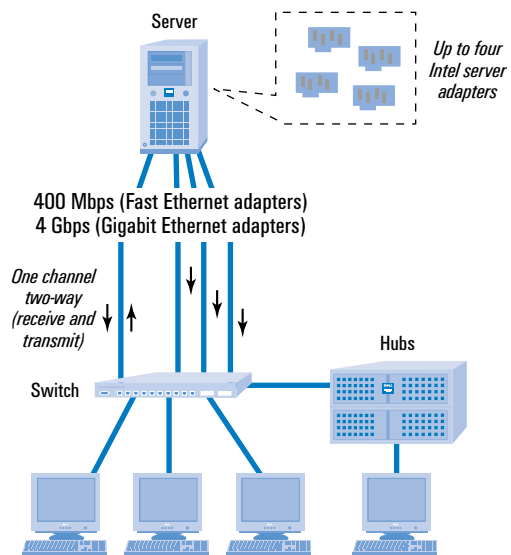


Figure 7. Balancing network traffic with Intel ALB




Figure 8. IEEE 802.3ad link aggregation

across multiple platforms and operating systems. Intel systems management tools include LANDesk Management, Wired for Management, Simple Network Management Protocol (SNMP) agents, and Desktop Management Interface (DMI). More recently, Intel authored and drove the Alerting Specification Forum (ASF) through the Distributed Management Task Force (DMTF). ASF provides the alerting infrastructure for desktops and servers, even in a powered-down state (see Figure 9). The alerts travel across the network to a management console, indicating events such as a contaminated desktop system, a pulled-out processor, or a failed hard drive.

Management tools and agents such as ASF and SNMP are currently available in Linux. Common Information Model (CIM), a model for describing overall management information in a network or enterprise environment, will soon be added to Linux.

In addition, consoles such as the Dell OpenManage systems management software can control and gather information about all of the components in the network, including storage limits, network connectivity, performance issues, downed servers, chassis intrusion, and host system memory parity errors. Intel hardware and software help enable this functionality.

The performance gains offered by TSO were so substantial in the Microsoft operating systems that Intel took advantage of TSO in the Linux environment.

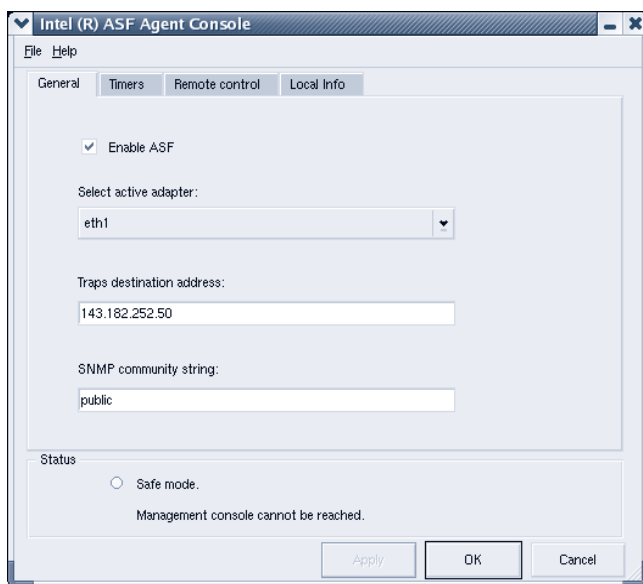


Figure 9. Monitoring alerts with the ASF agent console



Figure 10. Managing servers with OpenManage Server Administrator

**OpenManage Server Administrator.** This secure browser-based tool helps administrators manage individual servers. Server Administrator enables administrators with an Internet or intranet connection to manage the server from any location using a standard Web browser (see Figure 10). Administrators can obtain status information, troubleshoot, configure, and update systems management tasks. Integrated security is provided using Secure Sockets Layer (SSL) encryption.

**OpenManage Array Manager.** This online storage management application allows administrators to configure and manage local and remote storage devices attached to a server without bringing down the server. Array Manager provides a single console for array and volume configuration.

### Continuing support for Linux

The Linux operating system can help to provide the reliability, availability, scalability, and manageability needed in enterprise computing environments. Running on Dell PowerEdge servers configured with Intel PRO network connections, Linux offers IT managers a cost-effective operating system that does not sacrifice features or functionality.

**Gary Gumanow** (gary.gumanow@intel.com) is a senior technical marketing engineer in the Platform Networking Group at Intel Corporation. He is responsible for product planning for Intel Gigabit Ethernet silicon, adapters, and software. Gary has an M.B.A. from Pace University.

FOR MORE INFORMATION

<http://www.intel.com/network/connectivity>