

INTELLIGENT I/O

With the Internet growing at tremendous speeds plus the wide spread use of company intranets the demand for higher performance servers has skyrocketed. Host processors such as Intel's Pentium® or Pentium® Pro processors were designed to meet this high-performance processing demand. These processors integrate such features as 14- to 16-stage pipelines, out of order execution, and transaction-oriented buses to improve performance. In order to optimize these features the CPU demands higher-bandwidth input/output (I/O) to keep the high-speed host processor fed with data.

New network and storage protocols, such as Fast Ethernet and ATM, can provide high-bandwidth connectivity. However increasing the I/O bandwidth by an order of magnitude causes an increase in the number of low-level I/O interrupts sent to the host CPU. The solution? Allow the applications to be managed by a high-performance processor, like the Pentium Pro processor, and implement intelligent I/O for I/O processing.

What is Intelligent I/O?

When you hear the phrase 'Intelligent I/O' it most commonly refers to any server that includes a processor as part of the I/O subsystem. The I/O processor can perform tasks that would normally be executed by the host CPU. The concept is to reduce the overhead on the host CPU by off-loading interrupts and I/O processing functions, melding them into I/O subsystems managed by independent processors. This can greatly increase I/O throughput and decrease overall system response time.

Why Intelligent I/O?

Today's computing environment is changing rapidly along many dimensions. The increased performance levels of today's host CPUs highlights one major area in need of innovation - data I/O paths to and from the servers in the client/server model. The following factors require innovative I/O designs:

- Networked computing usage is replacing the stand-alone computing model, which in turn drives the need for network computing I/O.
- Networked computers increasingly aggregate vast quantities of data at the powerful server systems.
- Since the host processors (CPUs) in these servers now also run user applications, they need more powerful

storage interfaces for accessing larger and larger disk storage areas.

- Simultaneously, the sizes of files and messages are exploding as users move from simple ASCII data messages to complex, typeset-quality documents. These documents increasingly contain natural data elements that require huge file sizes.

THE I₂O* INITIATIVE.

The factors driving the switch to intelligent I/O systems have been powerful. OEMs increasingly are constructing servers with intelligent subsystems to meet the host CPU demands. But as that trend accelerates, designers are struggling with the software interface to the network operating system. The software interfaces need to be robust, abstract and standardized across all of the different operating systems and revisions.

Enter the I₂O Initiative

Early in 1996, a group of industry leaders formed a special interest group to define a standard interface for high-performance I/O systems. The resulting specification, dubbed I₂O, is aimed at vendors of systems, network and peripheral interface cards, and operating systems. The goals are to ease the task of building I/O systems, make it possible to distribute I/O functions across multiple processors, and dramatically boost I/O and overall system performance.

The resulting I₂O specification is based on an intelligent message-passing architecture and a split-driver model. The combination provides drivers that are portable across multiple operating systems and host platforms. System performance climbs because I/O and host operating system processing tasks can be distributed across multiple parallel processors.

How I₂O Helps Break the Throughput Bottleneck

To attain portability across multiple operating systems and host platforms, I₂O drivers are divided into two parts, or modules: The OS Services Module (OSM), and Hardware Device Module (HDM). The first module interfaces with the host operating system and the second interfaces with the particular device, media, or server managed by the driver. The two modules interface with each other through a two-layered communications system. A Message Layer sets up a communications session, and a Transport Layer defines how information will be shared; the Message Layer resides on the Transport Layer.



When combined with an execution environment and configuration interface, the I₂O communications model provides the HDM with a host-independent interface. Modules communicate without knowledge of underlying bus architectures or system topologies; messages take the form of a meta-language, so that communications do not depend on host operating system interfaces or bus topologies.

THE i960® MICROPROCESSOR POWERS I/O INNOVATION

With intelligent I/O established as a solution to relieving host CPUs from low-level I/O interrupts and the I₂O standard well on its way, the only missing element is an I/O processor to do the job.

The Intelligent I/O Processor

In a single chip, the i960 RP I/O processor brings all the basic elements needed to create an intelligent I/O subsystem in client/server networks and storage systems. An on-board PCI-to-PCI bridge enables designers to connect I/O components directly to the PCI bus and also add additional PCI slots to PC servers. In addition, the bridge improves overall system performance by reducing bus traffic and offers other features to provide flexibility in creating a complete PCI system.

In meeting the requirements for a high-performance system, the i960 RP I/O processor concentrates on five key areas: bus bandwidth, PCI availability, processing performance, data integrity and flexibility. Integrated around an i960 JF processor core are peripherals needed to build an intelligent I/O subsystem: the PCI to PCI bridge, an address translation unit for direct access between PCI and the local bus, a messaging unit, a DMA controller, a memory controller, and a PCI arbiter for the secondary PCI bus. Also included are two PCI buses, a local bus and DRAM bus, an I²C bus, and the APIC, the advanced programmable interrupt controller bus. Other on-board features help ease both hardware and software design.

Architecture Roadmap Importance

There is more to Intel's I/O processors than the processor itself. As I/O innovators consider applying 32-bit RISC power to their product plans, they must also consider the ability to track the performance of the host processor in the target server system. Because the i960 microprocessors are on the same technology track as Intel Architecture (IA) host CPUs, I/O innovators are assured that the architecture will have higher performance processor implementations for their subsequent product generations. Helping you onto this high speed

track is the diverse i960 microprocessor family, offering a wide breadth of price and performance. Intel offers an i960 microprocessor at a price suitable for almost any application.

The Right Building Blocks

In addition to outstanding service, diversity, and processor power the i960 processor family offers the most complete set of PCI building blocks for the new wave of innovative I/O products. With the Solutions960® Program offering over 100 different products and tools to support the i960 architecture there is no slowing down.

Intel offers a complete i960 Processor PCI Software Development Kit (SDK) specifically developed for I/O innovators. The software development kit contains everything developers need to begin I/O software development in a PCI system, including compilers, debugger, monitor and the base board.

To support the i960 RP I/O processor the IQ-SDK was developed. For ease-of-use the platform mimics the current Intel PCI-SDK. In addition the kit includes IxWorks*, an I₂O compatible RTOS developed by Wind River Systems. IxWorks provides a powerful real-time operating system that has clearly defined application program interfaces (APIs). All of this creates a user-friendly environment to rapidly develop basic device drivers, support NOS-to-driver independence, and allow multiple I/O software to coexist.

PUTTING THE PIECES TOGETHER

The explosive growth of the Internet and intranet, coupled with rapidly growing network and enterprise requirements, has generated a broad range of new client/server applications and products in the enterprise market: application, intranet, Internet, Web, video, firewall, and multifunction servers; data mining; data warehousing, SQL; transaction processing; groupware (i.e., Notes, e-mail); redundant disk arrays; high bandwidth networking; Fast Ethernet; asynchronous transfer mode (ATM); Fibre Channel; peer-to-peer; LAN-to-disk; inter-networking; routing; and switching.

Those diverse applications and products all have one common requirement: the need for robust I/O performance. Industry-leading computer vendors use intelligent I/O to maximize their system performance and to bring innovative and differentiable designs to market. The combination of the I₂O specification and a new intelligent-I/O processor architecture incorporated in Intel's i960 RP I/O processor now allows for the effective implementation of intelligent I/O.