Blade server – The booming server technology of the year 2002

By Steven Wu, Consultant

The rapid increase of Internet usage over recent years has severely clogged Internet traffic, creating a bottleneck for clients' online transactions. Also, with the upsurge of e-business galvanizing the marketing of most commercial sites, the IT industry has had to search for a solution to meet the requirements of availability and scalability for networking, computing, and storage. Traditional servers have had to adapt to meet the demand for a highly reliable and fast access to content for Internet users. This article explains how ultra-high density servers, or blade servers as named by the industry, can solve the headaches in server applications such as: space, clustering, redundancy, failover, and management problems.

Symmetric Multi-Processor (SMP) server and server clustering

The broad adaptation of the Symmetric Multi-Processor server, along with high performance Direct Attached Storage (DAS), were expected to meet the computing and storage demands of the Internet. However, scaling up the computing performance with a SMP server is very expensive. Server Clustering, an old but evolving technology, thus became a significant alternative to the SMP server. Thanks to the high bandwidth and low latency network technologies, Server Clustering interconnected numerous inexpensive servers (PCs) over a network allowing them to exchange simple messages to solve complex problems, just as SMP servers do. Generally, scaling up computing performance to Server Clustering is simple and inexpensive – connecting a server node (a PC) to the network forms a Server Clustering prototype. However, adding a well configured fail-over facility further allows Server Clustering to provide the high-availability service that an SMP server can never achieve.

Overcoming barriers to server clustering

In practice, Server Clustering is not only an alternative to an SMP server, but also a superior substitute for many applications. First, let's take a look at how different an SMP server box is from a Server Clustering. Most SMP servers reside in a tower enclosure; it is simple to manage such a box with plenty of room allocated for storage devices. Server Clustering, in contrast, could be a collection of several diverse PCs, network storage boxes, Ethernet switches, Ethernet/keyboard/mouse/display device cables, and power cords. With Server Clustering characterized in this way, cabling and space constrictions became a major management issue to its adopters. The arrival of a rack mounted solution sporting a uniform enclosure (19-inch 1U or 2U height) for the PC boxes two years ago solved this problem by allowing more PC boxes (42 server nodes for 1U, or 21 server nodes for 2U) to fit into a 19-inch 42U height chassis.

Evolution of the blade server

Recently, increasing the number of server nodes within the enclosure has materialized in the form of the blade server. The blade server (or High Density Server), molded each server node in a blade board form factor that no longer relies on its exclusive power supply unit and enclosure. Physically contained in a 19-inch rack enclosure, the blade server is a highly integrated powerhouse that vertically and horizontally stacks hundreds of servers. These high-density servers have eased the space limitation of 1U/2U server clustering. Capitalizing on this opportunity, and the company's technological expertise, NEXCOM announced its high-density server line HiServer in December 2000. Designed to run mission critical applications, HiServer with up to 20 server blades in a 3U height 19-inch rack mount chassis, is one of the most cost-effective server platforms available. Ultra-dense computing environments can now run more reliable applications and network services.

Built of high-grade steel, the HiServers can withstand the environmental conditions of the most rigid applications. Several blade (server) boards are available for the HiServer line, each with its own unique specifications to meet several server needs. Evolving around chipsets from Intel, VIA, and ServerWorks, each blade features above average memory support, 10/100/1000 Base-Tx Ethernet, add-on Fibre Channel ports, an I/O expansion capability via PMC or PCI, and most importantly a wide array of microprocessors from VIA C3, Celeron extending to Pentium III, and Tualatin CPUs. With a rack-mount platform, the HiServer's system architecture pushes high-density to the limits with a single set of peripheral devices (CD-ROM, FDD drive, keyboard, video display, and mouse) shared by all the systems within the rack. This revolutionary product likewise features hot-swap technology that simplifies replacement of server boards, power modules, and cooling fans, thus minimizing service time.

How does the blade server differ from 1U/2U server clusters?

A quick run down of the unique characteristics of the HiServer line demonstrates its advantages over traditional 1U/2U clusters. With the server cluster form factor redesigned on the HiServers, consolidation of resources and data has meant a step forward for server clustering. Moreover, the limited storage space within the HiServer, versus the massive storage capacities of 1U/2U clusters, has meant using storage sharing as a step-up solution. The HiServer's shared storage solution is elaborated on in the storage solutions section of this article. Despite storage sharing being a minor issue, this solution accompanied by the Gigabit and Fibre Channel ports on the HiServer still promises more computing speed, capacity, and functionality for various applications. Also, cable management on the HiServer travels via the same backplane within the chassis. Incorporating a centralized type of keyboard/VGA/mouse (KVM) switch, an entire HiServer rack needs only one main unit with a KVM controller to serve the keyboard, display, and mouse operations of the rest. This is quite impossible for 1U/2U server clusters because of cable management drawbacks. The following is a list of the advantages of the HiServer line:

- Ultra-dense server accommodating 9 or 20 blades in 3U and 126 blades or 280 blades in a 42U cabinet

- Cutting-edge KVM daisy chain technology requiring a maximum of only 15 cords on 280 systems for the power cords and KVM cables
- Cable-less CF-KVM switch technology sharing only one set of KVM (Keyboard, VGA, Mouse) and CF (CD-ROM, FDD) without any cables
- Expandable I/O via 32/64 bits PCI featured with additional high-speed I/O
- Easy to use system management monitoring the health status; power up/down, or restart of servers centralized remotely
- Easy installation and maintenance with hot-swappable components, eliminating system power-down during maintenance

Applications

As mentioned, the blade server boards deliver the performance and specifications needed for today's Internet applications. The blades aim for applications ranging from entry-level/front-end Web servers, firewalls, and mail servers to the tough bandwidth needs of video broadcasting. Meeting the demands of heavy Internet traffic, the blade servers can maximize system performance and scalability and can be used for:

- Colocation/server consolidation
- Content provisioning server
- Windows terminal server cluster
- Application server cluster
- High-performance computing

Why is the blade server suitable for today's Internet applications?

Colocation/server consolidation

Enterprise server consolidation and Internet data center

The blade server plays a very significant role in the co-location and sharing of data (see Figure 1). With the large amount of data stored in IDCs, a server needs flexibility in configuration, dynamic resource assignment, quick deployment of tasks, easy management and maintenance, and a space-saving architecture. This is not an easy task for traditional 1U servers to achieve. However, the HiServer satisfies all these requirements. There are a number of operating systems available to fulfill this task with the HiServer: Windows NT4, Windows 2000, Sun Solaris, and Linux, and FreeBSD. The combination of these operating systems and the HiServer assures high computing power in space-constricted areas.

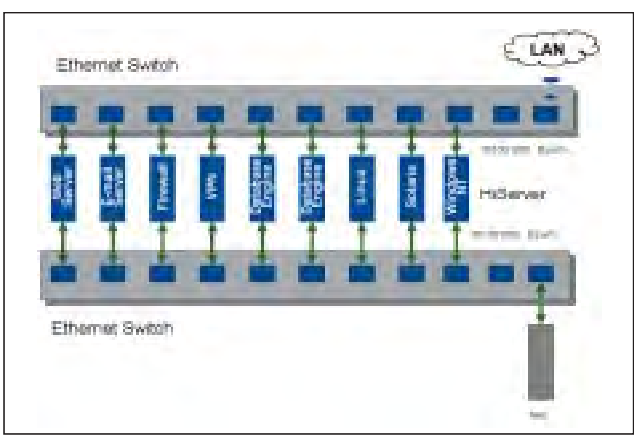

Figure 1

Load balanced terminal server

Web-centric operations and application service provider

Load balancing is a major factor in the network that squares off the indefinite number of requests to the server. Commercial Web sites normally deploy more than two Web servers to achieve load balancing. As requests bombard most Web-centric operations and application service provider facilities every second, these require the weighted sharing of all data transfers, processes, and communications within the network to avoid data flooding. Based on the illustration of the typical workflow for Web-centric operations and application service provider (see Figure 2), the use of a HiServer as a terminal server not only prevents any single device from being impaired, but also balances the server workload (I/O and computing speed), and scales the performance higher to accomplish the whole task. The HiServer eases management, maintenance, and quick deployment associated within the whole process. The HiServer functions as a load balancing server when combined with operating systems like Windows NT4/Terminal Server, Windows 2000/Terminal Server, Citrix MetaFrame, and VNC.

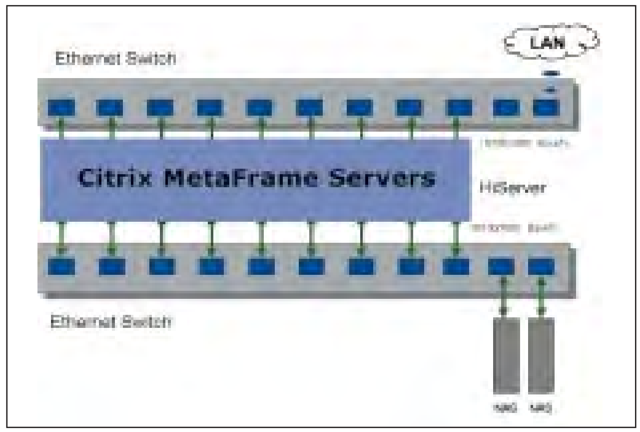

Figure 2

High-performance content provisioning server

Internet access, Internet portal site, Internet search engine, Internet game hosting

Oftentimes when we access the Internet, our workstations read data from or write data to a mass storage device. The speed with which we can locate a single byte of information on a Web site and download it for display is an important factor. From the xSP perspective, it is crucial to deploy a system that can serve the needs of its clients without having to suffer downtime. A typical example is game hosting. Another is an Internet portal site with operations such as e-mail, forums, search engines, and on-line shopping malls. All require a reliable system. The conversion of most search engines into portal sites, to attract more visitors, demands even higher computing power. Entry computing power, high disk I/O throughput, medium network bandwidth, intensive data access in file level, redundancy, automatic load balance, and dynamic configuration in a single system just about answer the requirements of these applications. Again, HiServer's flexibility allows portal and game hosting sites to realize these goals (see Figure 3). With operating systems such as Windows NT4 Server Cluster, Windows 2000 Server Cluster, Linux MOSIX/HA/NAT project, NAS/CIFS/Samba, IP take over/MAC take over, and Network Cache, these feats are as simple as counting 1-2-3 for the HiServer.

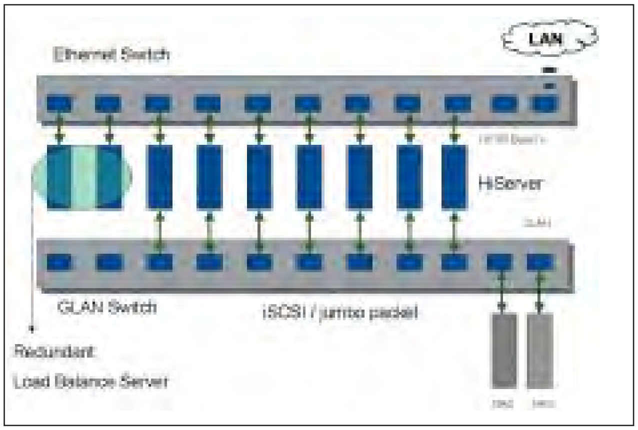

Figure 3

Application server cluster

Online transaction (processing, decision support system, broker information management, ERP, CRM, telco billing) and e-commerce

The search for high computing power, high disk I/O throughput, high network bandwidth, intensive data access, low message passing latency, and fast fail over are main concerns of the online transactions and e-commerce applications. The complexity of content within these network applications necessitates the HiServer to store and retrieve data in a parallel sequence. When the HiServer is applied on these infrastructures, the administrator can utilize the Application Programming Interface (API) and directly control the server without going through the DBMS's user interface. Once again the HiServer exhibits its strengths in completing tasks that manages thousands, even millions of users (see Figure 4). The HiServer as an Application Server runs well with operating systems such as Windows NT4 Server, Windows 2000 Server, Solaris, Linux, and J2EE application server solution.

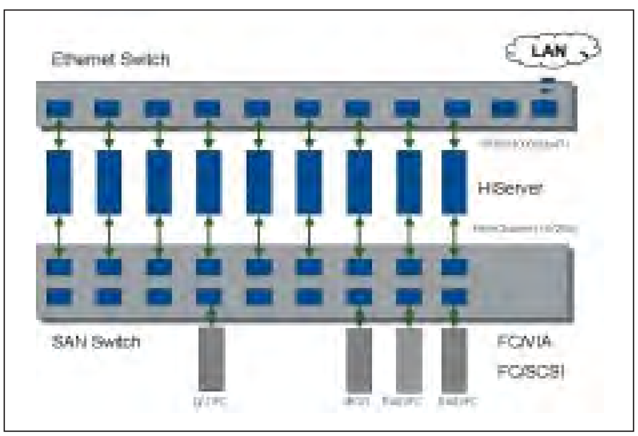

Figure 4

High-performance computing

Aerospace/astronomy research, weather forecast, genome analysis, chemistry/bioinformatics/aerodynamic modeling and simulation, and 3D rendering

High-performance computing focuses on developing supercomputers and software that run them. With the development of software programs that can be executed simultaneously by separate processors at its core, the HiServer permits and boosts parallel processing and triggers more computing power to the whole server structure without having to use any distributed processing software common with single-CPU systems. This is essential when acquiring information and processing data in industries requiring very high accuracy. Ultra-high computing power, ultra-low message passing latency, high network bandwidth, and medium disk I/O throughput features of the HiServer make it appropriate for high-performance computing (see Figure 5). The HiServer meets the demands of high-performance computing when used with Windows NT4 Server Cluster, Windows 2000 Server Cluster, Solaris Cluster, Linux Cluster, Beowulf, MOSIX, PPLINUX, SHRIMP project, PVM, MPI, VIA message passing protocol, and DQS cluster management.

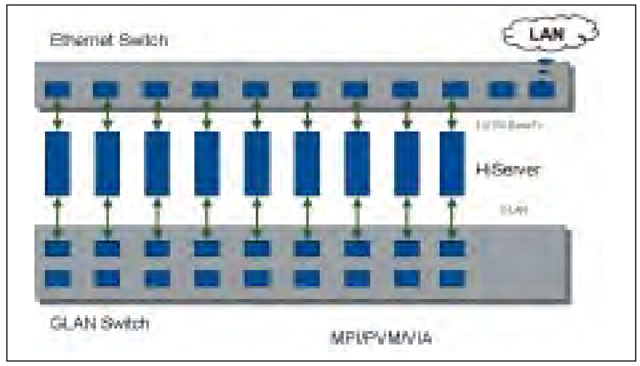

Figure 5

Storage solutions

Storage space is a major factor when shifting from 1U/2U server clusters into the HiServer architecture. As explained earlier, the HiServer utilizes storage sharing when greater storage capacities and computing power are needed. The remainder of this article describes the storage solutions applicable with the HiServer product line from NEXCOM.

SCSI connectivity

The structure of 1U/2U server clusters allows the presence of SCSI interconnectivity (hard drives) operated within one enclosure. Though this strongly supports its all-in-one attributes, there are drawbacks in terms of computing prowess and speed. However, the HiServer's shared storage is confined to the 2.5-inch drives attached on the blades. There is the possibility of utilizing a proprietary PCI Mezzanine Card (PMC), mounted on 2 blade servers, that can be interconnected to an external RAID storage. As illustrated in Figure 6, shared storage is coursed through blades 1 and 2 via the SCSI interfaced PMC card. With this structure, disk mirroring and data sharing/transfer exceeds the basic characteristics offered by 1U/2U server clusters.

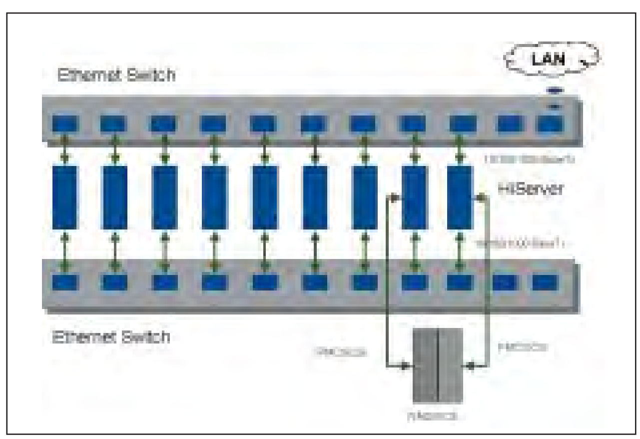

Figure 6

NAS Connectivity

NAS optimizes the traditional file server approach by creating systems designed specifically for data storage. Instead of starting with a general-purpose computer and configuring or removing features from that base, NAS designs begin with the bare-bones components necessary to support file transfers and add features "from the bottom up."

A NAS can store any data that appears in the form of files, such as e-mail boxes, Web content, remote system backups, and so on. Overall, the uses of a NAS parallel those of traditional file servers (see Figure 7). NAS systems strive for reliable operation and easy administration. They often include built-in features such as disk space quotas, secure authentication, or the automatic sending of e-mail alerts should an error be detected. The new breed of NAS networking products has succeeded in providing a reasonable alternative to traditional file servers in client/server networks. Besides its reasonable cost, a NAS promises reliable operation and easy management.

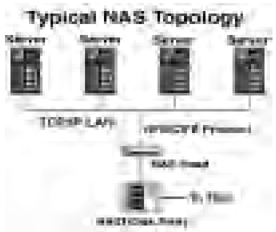

Figure 7

Evolving once again from its 10/100/1000 Mbits/sec Base-Tx Ethernet ports, the HiServer fulfills the dreams of NAS. It too provides an all-in-one NAS solution with all load balancing, servers, and NAS head operations contained inside it (see Figure 8).

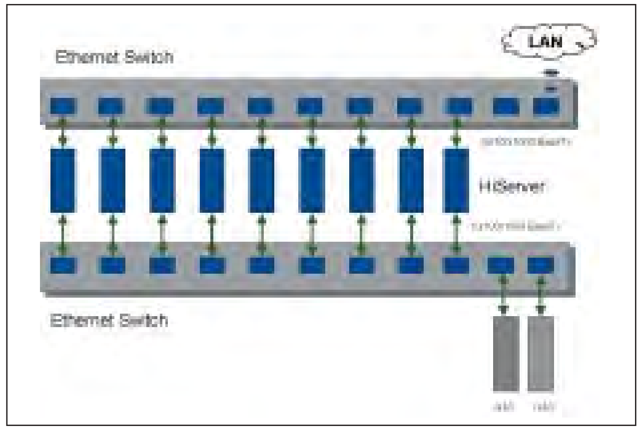

Figure 8

SAN Connectivity

The idea behind a SAN is radically different from NAS. To begin with, different protocols are used. In most SAN implementations, Fibre-channel (FC) adapters provide physical connectivity between servers, disk arrays, and tape libraries (see Figure 9). SAN represents a second network that supplements your existing

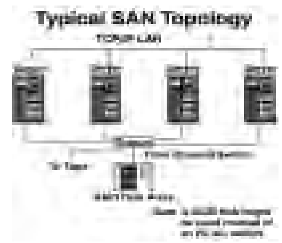

Figure 9

LAN. The advantage of a SAN is that SCSI is optimized for transferring large chunks of data across a reliable connection. Having a second network also off-loads much of the traffic from the LAN, freeing up capacity for other uses. Most significant is that backups no longer travel over the regular LAN and thus have no impact on network performance. How does the HiServer convert into a SAN?

Streamlined by the Fibre Channel ports on the HiServer, four of the blades (via a SAN switch) can connect to a DLT (Digital Linear Tape) of 20 to over 40 Gbytes and transfer rates of 2.5 Mbits/sec, a couple of RAID storage devices, and JBOD (Just a Bunch of Disks) that improves performance and fault tolerance of the system. Connection to the payload network is fulfilled using the 10/100/1000 Base-Tx Ethernet ports on each blade connected to an external Ethernet switch (See Figure 10).

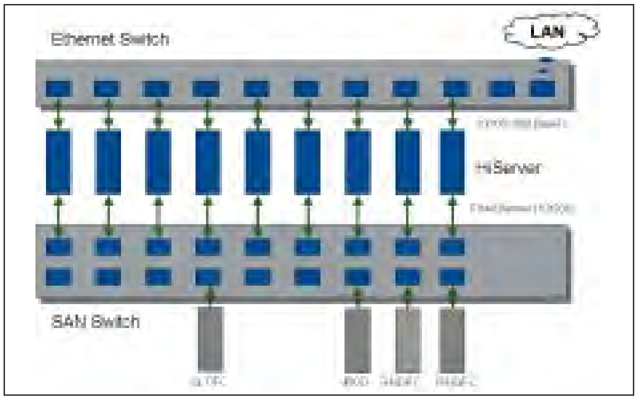

Figure 10

With such a structure, fault tolerance, network computing, and mirroring are pushed to the limits and stretches the capabilities of the blade server into new heights.

Referenced Sources

The Business Mac, Storage Solutions for your Network: SAN vs. NAS by Paul Shields

Introduction to NAS (Network Attached Storage) by Bradley Mitchell

***Steven Wu** has been Director of NEXCOM's Product Marketing Division, Server Business Unit since January 2001. Steven has also worked for a number of the top 10 IT manufacturers in Taiwan. These include Mitac, Arima, Clevo, and Elite. Steven has considerable experience in the IT field, having worked for six years in research and development with Symmetric Multiple Processor System; four years in software research and development for Notebook PCs; and three years in research and development and technical support for storage products, including DISK, Raid, and HDD.*

For more information, contact Steven at:

Steven Wu, Director, Product Marketing Division
NEXCOM International Co., LTD.
18F, No. 716, Chung-Cheng Rd.,
Cheng-Ho City, Taipei Hsien 235, Taiwan, R.O.C.
Tel: 886-2-82280500 (for English Service)
Fax: 886-2-82280506
E-mail: stevenwu@nexcom.com.tw
Web site: www.nexcom.com.tw