Server Blade Trade Association
Newsletter
June 10, 2003

Adoption of Blades Drives New Requirements for Storage

All Blades Are Not Created Equal

The concept of blades is often lumped into one large category when, in fact, there are several blade architectures that address specific market segment needs and offer unique value propositions. As the adoption of bladed computing architectures increases, there are profound implications for storage devices, storage systems, and the management layer between storage and servers.

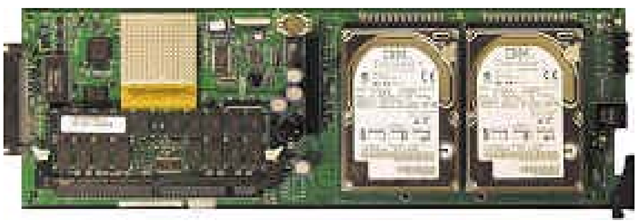

Web/Edge Server Blades

The first generation of blade servers were designed to address the problems associated with rapidly rising costs of data center and hosting provider environments. Optimization of space, power consumption, and heat generation were the primary metrics driving the value proposition of this class of server blades.

Typically, web/edge blade servers have 1-2 Pentium class processors (or 64bit Sparc), connect to the backplane through Ethernet, and have blade resident hard disk storage. The local hard disk is primarily a local cache storage that stores the application image, configuration, management system hooks, and other read only data. Small form factor, power consumption, and cost are the key requirements for this class of blade server. These requirements are best served by hard disks with ATA, SATA (Serial ATA), and SAS (Serial Attached SCSI) interfaces. Although current blade server products use 3.5-iinch and 2.5-inch disk drives, because capacity is not the dominant driver, 1.8-inch or one-inch form factor hard disks may eventually find opportunity in this market.

Flexible Midrange Server Blades

Flexible Midrange server blades were developed to address customer needs to easily configure multiple SMP servers to meet current needs with the flexibility to adapt to future requirements without a forklift upgrade. Additionally, the need for modular scalability is a key feature of this class of blade server. Optimization of a switched backplane for performance and high availability are the primary metrics driving the value proposition of this class of server blades.

Typically, flexible midrange blade servers have 1-4 Xeon class processors (future 64bit Intel architecture and Sparc), connect to the backplane through Infiniband (or a proprietary switched backplane that is similar to Infiniband), and do not have blade resident hard disk storage. Storage is accessed through an I/O controller blade that contains both Ethernet and Fiber Channel ports for access to NAS and SAN storage subsystems respectively.

Network Attached Storage (NAS) and Storage Area Network (SAN) subsystems (which may themselves contain storage blades) will drive requirements for hard disks serving this class of blade server. High performance SCSI/SAS, Fiber Channel, and SATA hard disk drives with 4-5 platters will find opportunity in this market. Volumetric storage capacity and system level performance will drive hard disk requirements for this market. SATA will see significant opportunity in this market as the performance of many SATA drives with RAID rivals traditional high end hard disk drives at a fraction of the cost disk drives with SCSI and Fiber Channel interfaces.

Storage Blades

In addition to server blades, a few vendors are introducing products that can be classified as storage blades. Typically, this class of blade provides near-line data protection and caching services with the benefit of fine grained scalability. This enables customers to pay as they grow and take advantage of the ever increasing increase in areal density.

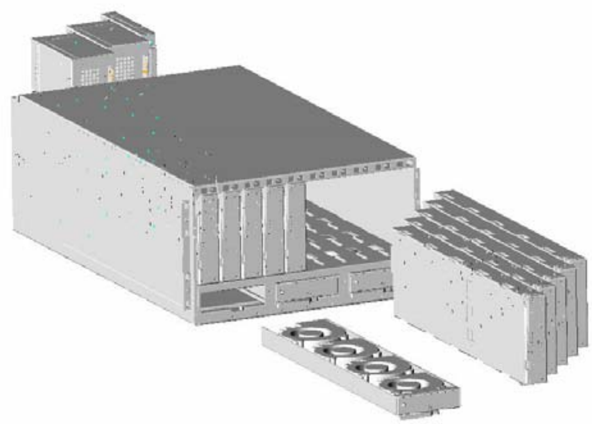

Current storage blade products utilize high capacity 3.5-inch SATA disk drives. This can provide a storage blade with over a Terabyte per blade. With the increase in the areal density of magnetic hard disk drives the 3.5-inch drive blades will be able to provide several Terabytes of storage by 2006. However a smaller form factor disk drive can be used to make smaller blades with capacities comparable to the current generation of storage blades. The figure below illustrates the effect of recording areal density growth relative to form factor for a given storage capacity.

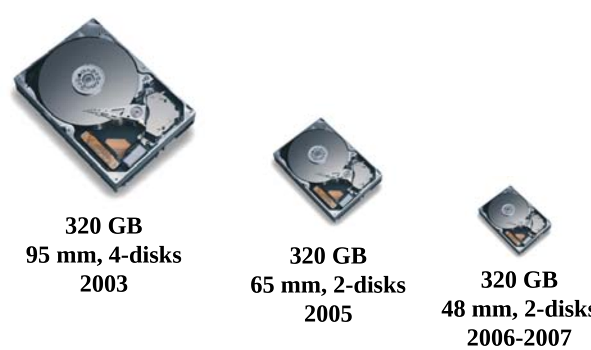

Storage Blades can be used to make NAS appliances with much higher aggregate capacity per removable storage module (the blade). The drives on a storage blade can also be organized using RAID architectures to make a virtual disk drive of the storage blade with very high capacities and data transfer rates. Storage Blades can also be part of a SAN storage system using switched networks with Fibre Channel or Ethernet network protocols. There has also been some work on switched fabrics using external SATA networks that could also be used to create a SAN network.

Rapid Increase in Areal Density Brings Opportunity

The recent developments in blade-based computer products are made possible by the rapid evolution in electronic integration and magnetic hard disk drive recording areal density (the amount of information that can be packed on a disk surface). Although the rate of areal density increase has been slowing gradually, two and a half years ago the cumulative annual areal density growth was well over 100% annually. This growth in recording areal density has led to very high capacity 3.5-inch hard disk drives (currently up to 320 GB with 4 platters) and very useable capacity 2.5-inch and smaller form factor hard disk drives. The figure below shows expected capacity growth of 2-platter 2.5-inch (65 mm), 1.8-inch (48 mm) and 1-inch (27 mm) disk drives and a reference point for a 3.5-inch (95 mm) 4-platter capacity in the 2005-2006 time frame.

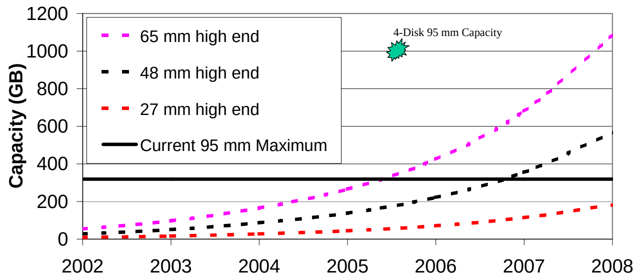

Another important trend in disk drives has been the development of very high data rate low cost drive interfaces such as SATA. The current generation SATA support 1.5 Gbps data rates while the next generation SATA, available next year, will double that available data rate. Since SAS interfaces that will provide SCSI command support on the same serial physical interface as SATA will also be available next year there may be some interesting combinations possible for blade based storage systems that mix the cost advantages of SATA and the more comprehensive command support of the SCSI interface.

Challenges of Blades and Networked Storage

Blade server architectures that separate processing and storage do so to create a system that has independent scalability of server, networking, and storage resources. This blade architecture leverages the current trend towards networked storage whether that is NAS (Network Attached Storage) for file-level data access or SAN (Storage Area Network) for access to block-level data in database environments. Currently, access to SAN's is primarily through Fiber Channel networks but there is a strong financial case for blade servers that access SAN's through Ethernet networks using protocols such as iSCSI.

In environments where blade servers access SAN's, scalability of clustering mechanisms for failover and load balancing will be extremely important to the adoption of blade architectures in high availability environments. Additionally, storage management systems will need to adapt to environments where computing resources and applications are redeployed dynamically through automated management tools.

This dynamic redeployment plus the increase in the number of independent computing and storage resources make management crucial to the growth of complex blade based architectures.

Blade server environments that utilize NAS as their external storage will also require more sophisticated scalability, fail-over, and load balancing mechanisms for NFS and CIFS protocol environments. Provisioning, sharing and protection of data in a dynamic environment is a crucial and non-trivial challenge.

Many of the blade storage architectures utilize parallel or serial ATA disk drives. ATA disk drives and the drive interface were not initially designed for use in enterprise environments with 24X7 usage, nor were they designed with the sophisticated storage commands that have been developed for SCSI disk drives. Today's higher capacity ATA drives are designed more along the lines of traditional enterprise disk drives but it will be difficult to create the same reliability level without increasing the price of these drives. One approach to increase the reliability of these drives is through use of more sophisticated SMART drive algorithms. SMART is a self-monitoring facility built into ATA (as well as other interface) disk drives that can be used to predict drive failure. While conventional SMART algorithms are probably not more than 50% accurate CMRR at UCSD has developed an advance SMART algorithm that has been demonstrated to improve this predictability considerably.

SMART is one example of added intelligence on disk drives to improve drive performance. It may be that use of added drive intelligence features in ATA disk drives, I-Drives, could be a key to better management of dense storage environments such as an array of storage blades in a NAS or SAN configuration. Some local file attributes or even object-based storage concepts could assist with some drive functions and storage management. On-drive encryption, data searches, and fast secure erase may add extra functionality or security to the data in these storage arrays. The increasing level of electronic integration and developments in drive firmware could play a key role in the simplification of storage management that is becoming an increasing issue in the proliferation of storage networks.

Conclusion

Vendors that deliver the best value to customers by recognizing the new requirements of the respective blade market segments, regardless of the vendor's place in the value chain, will reap the largest rewards. Data storage in the blade environment will include the original application in web/edge servers, flexible midrange server blades without much local storage and storage intensive storage blades in NAS appliances or as part of a SAN architecture.

The growth in data storage areal density will make very high volumetric blade based storage systems possible. With higher storage density, single 3.5-inch drives will be available with capacities over a Terabyte or smaller form factor drives will be available with 100's of Gigabytes in just a few years time. This allows significant increase in local cache on a web/edge server blade in a smaller footprint and makes blade storage systems possible that are much smaller or that have many more resident hard disk drives than the current generation of products.

The growth of the blade architecture and in particular the growth in future storage blade architectures will likely be an important driver in data storage products. This will also be a vehicle for greater proliferation of SATA-based storage products. If this trend is combined with use of lower cost Fibre Channel or SCSI-based Ethernet storage networking, the result will be a continued trend in more affordable data storage systems.

It is likely that blade storage growth will be limited by the capability of storage management whether for NAS or SAN storage environments. A key requirement for the proliferation of storage blade architectures will be the storage management infrastructure to support it. It is possible that blade based storage will benefit from greater on-board drive intelligence.

Authors:

Thomas Coughlin, Coughlin Associates tom@tomcoughlin.com
Patrick Hanlon, Independent Analyst and Consultant phanlon@mindspring.com