

THE COMPETITIVE ENVIRONMENT FOR STORAGE AREA NETWORKING

Until the late 1990s, external data storage was entirely DAS, or directly attached storage, that is, storage attached directly to an application server. Even today, more than half of all external storage is DAS. Increasingly, storage is moving toward being attached through a network, either as NAS (network attached storage) a file-based system connected to application servers through a LAN, or SAN (storage area network) a block-based system connecting application servers to storage through a high-speed dedicated network using switches and hubs.

Because the storage equipment market is so large, \$40 billion, the variety of companies and the product offerings is huge. \$40 billion in hardware, software, and service sales means that even crumbs from the major competitors' tables can provide a nice business for many small niche players. Or not.*

Competitors range in size from IBM to start-ups with \$250,000 in annual sales. Products range from simple software applications and stand-alone storage boxes to end-to-end SAN systems. In this section we will somewhat oversimplify that complex environment to provide some perspective on the type and quality of organizations and technologies in this market space and the reasons why we believe we can win market share.

Within the SAN/NAS market competitors offer five basic types of storage products:

- End-to-end solutions, such as Sun's *N1*.
- Server-based software solutions, such as Veritas or Tivoli's Storage Resource Managers
- Storage-based solutions, such as Network Appliance's NAS or EMC's *Symmetrix*
- Network-based solutions, such as H-P's *StorageApps*, or Brocade's *Rhapsody*
- Application software, such as FalconStor's or Kom Network's virtualization software.

END-TO-END SOLUTIONS

The storage market has been, and still is, dominated by this type of competitor --the big vertically integrated OEMs such as IBM, Sun, and H-P who offer total solutions from computing through storage. Until the 1990s, the big OEMs shared the market among themselves by catering to captive customers, only occasionally bothered by plug-compatible vendors such as Memorex who competed on price. But storage never has been about price; it's about reliability and customer support. Until EMC came along to beat the OEMs at their own game, the big integrated OEMs were slow to innovate in storage.

EMC showed that customers would pay more for innovative storage with new features and functions, such as Point-in-Time copying and a high level of service. EMC established its credibility first in mainframe storage, and in doing so, took substantial market share from IBM's largest customers. Once EMC proved it could deliver and support reliability, EMC quickly went on to take market share in the Open Storage market segment against all three of the vertically integrated OEMs. H-P paid EMC more than \$500 million a year for storage it resold to its own customers.

EMC sales and earnings peaked in 2000, with sales of \$9 billion and a gross margin of 64%. IBM, Sun, and H-P began gaining back market share in 2001 and 2002 while EMC sales and earnings started heading south. EMC sales in 2002 were down 40% compared to 2000. IBM trumpeted its increased storage sales in 2001 despite a softening storage market.

* Cereva, a \$160 million startup folded last June. Cereva was developing a large storage array to compete with EMC's *Symmetrix*, IBM's *Shark*, and Hitachi's *Thunder* and *Lighting* arrays, a market segment with shrinking margins and little product differentiation.

Despite IBM's little comeback, starting in 2002, it became clear that storage had become a commodity, price pressures are severe, and IT end users are placing little value on vendors' networks and new features. Virtualization, the big hype of 2000, has fallen flat because of interoperability problems.

But those OEM end-to-end systems do not interoperate with those of other vendors, nor with many third party applications, and that is increasingly disturbing to the IT customers. The IT industry wants standards. The OEMs like proprietary, captive customers, so they offer only token standardization.

SERVER-BASED SOLUTIONS

Many large and small companies are offering storage management and virtualization solutions that run on storage servers. Two of the largest and most successful are Veritas and Tivoli, the latter a subsidiary of IBM. Although server-based solutions are able to virtualize storage across arrays from many different vendors, they do not work well when the network incorporates hosts from many vendors with different operating systems, so they tend to be used only with homogeneous, end-to-end systems.

Both Veritas and Tivoli consider data protection, application availability, and management their core capabilities, rather than virtualization and consolidation. They tend to partner with manufacturers of equipment at the perimeter of the network rather than developing interface adapters themselves. That limits interoperability to partners' equipment and homogeneous hosts.

Server-based systems could run on the processors which make up our Storage Operating Environment, as they also can run on director-level switches. In this sense, we are a potential partner, or, as Veritas terms it, we would be Veritas Enabled, and it would permit Veritas to interface more easily with heterogeneous hosts.

Veritas is a strong supporter of industry standards, such as the Common Interface Model (CIM) being proposed by the Storage Networking Industry Association because it would make server-based systems much easier to implement.

STORAGE-BASED SOLUTIONS

Manufacturers of storage-based solutions, such as EMC, claim that they can virtualize and manage storage, which is true so long as all of the storage is homogeneous and, often, contained within the same physical cabinet. Storage-based solutions use an insular file system, which means that for an application to find a file it must be mapped to the right cabinet. Because such systems are not easily scalable, when storage demands reach multi-terabyte levels, storage-based solutions tend to be limited to corporate Data Centers, to workgroups, or to small and medium-sized organizations.

EMC pioneered in allowing multiple servers from different manufacturers to access storage, and in providing point-in-time copy services. As the intelligence in EMC's systems increased, EMC also became Veritas Enabled, so that storage management software could be run on an EMC system. But as the major end-to-end OEMs caught up with EMC with intelligent storage features, EMC's market share began to fall.

NAS systems are a special case of storage-based solutions, with their own limitations. Although they are less costly than SAN-type storage, they also are lower in performance and are poor solutions for back-up and recovery. NAS systems are file-based, rather than block-based, as SAN storage is, so it is easier for users to share data. That advantage appeals to workgroups, too.

NETWORK-BASED SOLUTIONS

There is a strong consensus developing in the IT industry that network-based solutions are the best way to consolidate and virtualize storage, and to provide application services. But it is a complex issue, for there are at least four different architectures vying for the lead.

1. *Symmetric appliances*, such as StorageApps and DataCore, are intelligent workstations that are in the data stream, so that the data storage stream and the data virtualization and management instructions share the same path, which can create a bottleneck. A symmetric SAN appliance can only virtualize storage to which it is directly attached. Multiple appliances form islands of storage, not a virtual pool across the entire SAN, which limits scalability.
2. *Asymmetric appliances*, such as StoreAge and VersaStor, are intelligent workstations in which the storage data and virtualization and management instructions use separate paths. This type of appliance may perform better than a symmetric appliance in data transfer. Both symmetric and asymmetric appliances require an agent on each host to direct I/O commands to the appliance. Use of a single workstation to control a SAN presents both a reliability. Because all appliances proposed to date are Wintel machines, the reliability issue is amplified. Dual workstations are only a slight improvement. If appliances are used to virtualize JBODs, then lower-performance RAID functionality must be performed by the host.
3. *Switch-based solutions* are switches with enhanced intelligence and an operating system to run virtualization and storage management software, which software could be independent of the switch manufacturer. One advantage of a switch-based solution is that the switch memory could be used to provide inexpensive caching. But a disadvantage is the introduction of a second operating system which may not be robust enough. Dual switches improve robustness, but increase cost dramatically. Moreover, a switch is essentially a gate-keeper, which involves a delay that degrades switch performance.
4. *Router-based solutions*, such as offered by Vicom Systems, are intelligent devices that select the path data takes through a network, and translate one protocol into another in the process. The routers also have enough intelligence to host storage virtualization software and applications such as copy services, data migration, RAID, and remote back-up. The router solution is scalable and self-healing with no single point of failure.

All of the four types of network-based solutions present the IT administrator with a GUI to manage the SAN and to pool and allocate storage. All of them can set up zones of storage which can be governed by policy, either with their own or third-party software.

APPLICATION SOFTWARE

Dozens of companies -- perhaps hundreds -- both small start-ups and large computer systems houses, have developed software applications for virtualizing storage, and for performing many of the tasks associated with it, such as file and protocol translation, security, copy services, remote back-up and restore, and data migration. The software is written to be run on servers, intelligent switches, and storage equipment.

None of those companies has been successful in selling virtualization software because virtualization has no benefit to IT managers unless the two barriers to consolidation can be overcome: interoperability and workgroup resistance.